

REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-01-

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be penalized for providing information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE AL

0178

1. REPORT DATE (DD-MM-YYYY) 01-01-2001			2. REPORT TYPE Final		3. DATES COVERED (From — To) 1 January 1998 — 31 December 2000	
4. TITLE AND SUBTITLE Control of Nonlinear Systems					5a. CONTRACT NUMBER F49620-98-1-0242	
					5b. GRANT NUMBER	
					5c. PROGRAM ELEMENT NUMBER 61102F	
					5d. PROJECT NUMBER 2304	
6. AUTHOR(S) Eduardo D. Sontag					5e. TASK NUMBER	
					5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Mathematics Rutgers, The State University Piscataway, NJ 08854-8019					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM, Building 410 Bolling AFB, DC 20332-6448					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approval for public release; distribution is unlimited.					<p>AIR FORCE OFFICE OF SCIENTIFIC RESEARCH (AFOSR) NOTICE OF TRANSMITTAL DTIC. THIS TECHNICAL REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLIC RELEASE LAW AFT 130-12. DISTRIBUTION IS UNLIMITED.</p>	
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the authors and should not be construed as an official U.S. Government position, policy or decision, unless so designated by other documentation.						
14. ABSTRACT This Final Report summarizes accomplishments of the grant research. The work reported here deals with the mathematical foundations of nonlinear systems analysis and feedback control. It describes advances in input to state stability, which is a notion of stability with respect to input signals, IOSS, which a notion of well-posed detectability, and related concepts, resulting in necessary and sufficient characterizations in terms of dissipation inequalities. Other topics include control-Lyapunov functions, the design of feedback controllers by means of nonsmooth analysis techniques including viscosity subgradients, and the introduction of new notions of stability with respect to finite-energy measures as well as their complete characterization.						
15. SUBJECT TERMS keywords; control, nonlinear systems, feedback, input to state stability						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT		18. NUMBER OF PAGES	
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	UU		29	
19a. NAME OF RESPONSIBLE PERSON Dr. Marc Jacobs					19b. TELEPHONE NUMBER (include area code) (202) 767-5027	

CONTROL OF NONLINEAR SYSTEMS

F49620-98-1-0242, Final Technical Report, 1/1/98-12/31/00

Eduardo D. Sontag
SYCON – Rutgers Center for Systems and Control
Department of Mathematics, Rutgers University
New Brunswick, NJ 08903

ABSTRACT

The work reported here deals with the mathematical foundations of nonlinear systems analysis and feedback control. It describes advances in input to state stability, which is a notion of stability with respect to input signals, IOSS, which is a notion of well-posed detectability, and related concepts, resulting in necessary and sufficient characterizations in terms of dissipation inequalities. Other topics include control-Lyapunov functions, the design of feedback controllers by means of nonsmooth analysis techniques including viscosity subgradients, and the introduction of new notions of stability with respect to finite-energy measures as well as their complete characterization.

20010326 136

1 Introduction

The work reported here deals with the *mathematical foundations of nonlinear systems analysis and feedback control*. The control of highly nonlinear systems is crucial to the Air Force mission, as nonlinear effects cannot be ignored in the design and modeling of high-performance aircraft.


Our emphasis is on the development of *basic principles*, and on the communication of these results to those in the engineering community (and, in particular, to other AFOSR researchers) who, in turn, employ these techniques in applications. As evidence of our track record in such knowledge transfer and of its value to other investigators, we may cite the following examples: (a) over 50 papers in the 1999 CDC Proceedings refer to the PI's work; (b) Petar Kokotovic's plenary address "Constructive nonlinear control: progress in the 90's" to the latest IFAC World Congress based a substantial portion of its exposition on concepts originally introduced and developed by the PI; (c) several recent textbooks and research monographs rely heavily upon the PI's work (e.g., the second volume of Isidori's text, as well as the latest edition of Khalil's popular nonlinear systems textbook, as well as books by Krstic et.al., Sepulchre et.al., Freeman et.al., etc.; (d) the PI is regularly invited to deliver plenary talks describing his work (e.g., 1999 British Applied Math conference, 1999 Nonlinear Analysis and Control workshop in Portugal, 2000 European Community Nonlinear Control meeting, course in Nonlinear Control at 1999 AMS Annual Meeting).

Feedback design technology is undergoing an exceptionally rich period of progress and maturation, fueled to a great extent by (1) the discovery of certain basic conceptual notions, and (2) the identification of classes of systems for which systematic decomposition approaches can result in effective and easily computable control laws. These two aspects are complementary, since the latter approaches are, typically, based upon the inductive verification of the validity of the former system properties under compositions. It is perhaps in the first of these aspects, and in particular in the precise formulation of questions of robustness with respect to disturbances, and of stabilization conditions, that the PI's previous work has had most influence, although our research covers a wide spectrum of other subjects as well.

1.1 ISS-Related Notions

A common feature of a large part of the project is the study of "sensitivities" of system quantities $y(t)$ to signals $d(t)$. The inputs d might represent errors, disturbances, or tracking signals, and the outputs y could stand for the entire state, or a more general quantity such as a tracking or regulation error, or the distance to a target set of states such as a desired periodic orbit.

Often, the symbolic diagram $d \rightarrow \square \rightarrow y$ represents a closed-loop plant (P) / controller (C) configuration, and d might decompose into a pair $(d_{\text{act}}, d_{\text{obs}})$ where d_{act} corresponds to actuator errors and disturbances, while d_{obs} summarizes measurement or observer errors or disturbances.



A central issue in this context concerns the appropriate formalization of "stability" of $d(\cdot) \mapsto y(\cdot)$, and, associated to this, the design of controllers which achieve such stability. There are two *fundamentally different* ways to interpret the effect of external signals d , depending on whether we wish to think of such signals d as "large" or "small". In the latter case, the goal is often to understand behavior which is insensitive to d (e.g., asymptotic disturbance rejection, robustness of stability in the presence of small observation errors and vanishing perturbations

on dynamics, no chattering due to noise), and this leads to the questions discussed briefly in the introductory Section 1.2 and in somewhat more detail in Section 2.1. On the other hand, for large d , interesting questions concern the characterization of the transient and steady state behaviors attributable to d . This leads to *input to state stability (ISS)*, introduced in the PI's 1989 coprime factorizations paper, which already appears as a foundational concept in several texts for advanced nonlinear control courses, as cited earlier, and continues to have a major impact on current research. As generalizations of "system gain," ISS and its cousins (IOS, etc.) afford a coherent and elegant qualitative measure of the mutual dependence of signals (inputs, states, disturbances, output objectives, measurements, etc), allowing the precise formulation of questions of robustness with respect to disturbances and of the well-posedness of identification and observation. We next informally overview the area, setting the stage for the questions to be discussed in more detail in Section 2.2.

Let us briefly recall the operator-theoretic approach to stability, as developed in landmark work by Zames, Sandberg, Desoer, Safanov, Vidyasagar, and others, starting in the 1960s, and including important recent extensions by Georgiou, Smith, and others. In this approach, typically, a "system" is a causal operator F between spaces of signals, and stability is taken to mean that F maps bounded inputs into bounded outputs, or finite-energy inputs into finite-energy outputs. More stringent requirements in this context are that the gain of F be finite (in more classical mathematical terms, that the operator be bounded), or that it have finite incremental gain (mathematically, that it be globally Lipschitz). This has been extended sometimes, especially in work of Safanov and Desoer, to allowing various types of nonlinear gain, or equivalent geometric properties for graphs. The input/output approach has been extremely successful in the robustness analysis of linear systems subject to nonlinear feedback and mild nonlinear uncertainties, and in general in the area that revolves around the various classical versions of the small-gain theorem. Moreover, geometric characterizations of robustness (gap metric and the like) are elegantly carried out in this framework, which also provides a natural setting in which to study the classification and parameterization of dynamic controllers.

A complementary aspect addressed by the "ISS" philosophy is the characterization of the effect of initial states on transient behavior, in a manner consistent with Lyapunov stability and dynamical systems thinking. The inclusion of terms involving initial conditions (or, in a non-state space formalism, the effect of past inputs) constitutes one of the main differences between the "ISS" approach and operator-theoretic definitions of stability; another is the systematic use of nonlinear gains and the emphasis on the study of equivalent conditions for existence of estimates of the various types.

One manner to introduce notions of the ISS type is by means of coordinate changes, starting from more classical concepts. In typical operator-theoretic stability analysis, one studies stability of input/state (or input/output) operators with respect to various induced norms: $L^2 \rightarrow L^2$ (H_∞ stability), $L^2 \rightarrow L^\infty$ (H_2 stability), or $L^\infty \rightarrow L^\infty$ (L^1 stability). In the context of the current discussion, we do not focus upon the actual numerical values of gains (operator norms) but, rather, on the more qualitative question of *existence* of estimates (finiteness of the gain, boundedness of operators with respect to the different norms); for finite-dimensional linear systems, of course, all these questions are equivalent, and they amount to the requirement that the origin should be asymptotically stable with respect to the unforced dynamics. Furthermore, we adopt in our approach the general principle that *notions of stability should be invariant under (nonlinear) changes of variables*. That is, we take what might be termed a "topological" (as opposed to a "metric") approach to the qualification of stability. Consider any differentiable manifold X with a distinguished point x^0 (in our application, X is the state space \mathbb{R}^n , or the

spaces of input or output values, and $x^0 = 0$ if we are interested in stability with respect to the origin and zero signals, but we can define matters in more abstract generality). By a *change of variables* in X , we shall mean a transformation $z = T(x)$, where $T : X \rightarrow X$ is a homeomorphism such that $T(x^0) = x^0$ whose restriction $T|_{X \setminus \{x^0\}}$ is a diffeomorphism (i.e., the only potential loss of regularity is at the point x^0).

Changes of state and input variables transform $L^\infty \rightarrow L^\infty$ gains for input/state operators immediately into the definition of ISS (for an exposition which emphasizes this point of view, see the recent survey [24], cf. the PI's web site, and Section 2.2). It is far less obvious that $L^2 \rightarrow L^2$ stability *also* leads, under coordinate changes, to ISS. This fact has now been precisely formulated and established in [4]; its proof relied upon the rather nontrivial characterizations of ISS obtained in 1996 work with Wang and it settled the question of interpreting ISS as a generalization of " H_∞ " stability. Another breakthrough that took place during the last couple of years concerns the notion that obtains from $L^2 \rightarrow L^\infty$ finite gain under coordinate changes: *integral* input to state stability (*iISS*), which is strictly weaker than ISS and can be interpreted as a form of finite-energy response. The recent fundamental work [9] provided several extremely elegant and natural alternative characterizations of iISS, for instance in terms of the existence of some output which makes the system 0-detectable and dissipative, and showed the relevance of iISS to tracking problems with "nonlinear resonance" behavior. Even newer research in this most active area, cf. [16], provided generalizations of "mixed norm" stabilities.

Other notable breakthroughs in this area within the last year or so include (1) the complete characterizations of input-to-output stability (*IOS*), see [26], [12], [14], which is a notion for systems with outputs which is geared towards the formulation of regulation problems, and (2) the completion of a program of research to show that detectability of nonlinear systems is equivalent to a dissipation condition with smooth storage functions (a problem on in which we had been working for several years, finally settled in our student Krichman's thesis dated January 2000), see [2], [40] and [19] (preprints available at the PI's web site).

In this context, it is worth mentioning a recent mathematically intriguing result from [15]. It was shown there that, conversely, it is possible to recover systems with finite " H_∞ gain" by a change of variables from ISS systems, provided that the state space has dimension different from 4 and 5. (For those particular dimensions, the problem remains open, since our proof relies heavily upon Smale and Milnor's results on h -cobordism, valid only in spaces of large dimensions.) The result implies, as a most special case, that global asymptotic stability of an equilibrium is equivalent, under a change of variables, to global *exponential* stability. (Obviously, this result does not say that Lyapunov's concept of stability is uninteresting because it is merely a distortion of exponential stability, nor that ISS is just H_∞ theory in disguise, since actually computing – as opposed to merely showing existence of – the appropriate changes of variables is not practical.) Since it provides further insight into the meaning of ISS, we would like to explore further implications of this result, as well as extensions to iISS.

In the (by now) classical theory of finite-dimensional linear systems, there are three basic notions of stability type: (1) internal stability (behavior of states), (2) detectability (external behavior determines how states evolve), and (3) input/output stability (transfer function has no poles in the right-hand plane). In perfect analogy – or rather, generalization – the three basic notions of stability: (1) ISS, (2) IOSS, and (3) IOS play the exact same role as the respective linear-system notions. For example, the theorem "a system is internally stable if and only if it is i/o stable and detectable" generalizes to "ISS equals IOSS and IOS". The notions of ISS, IOS, and IOSS (but called "detectability" and expressed in i/o terms only) were introduced by the PI in two 1989 papers. The triad consisting of ISS, IOS, and IOSS provides

a framework for the analysis of nonlinear systems; the complete characterization of all these properties in dissipation (Lyapunov or energy-like) terms has only been finalized during the last few months. To emphasize, the goal of the PI's work in this area is highly ambitious: the complete reformulation of the foundations of nonlinear control based on ISS-like ideas. This goal is very long-term, but definitely worth the effort: the payoff will be in the development of a consistent, elegant, and ultimately design-oriented, systematic approach to the subject. The work here reported supported research which led to truly exciting advances during the last three years, and we hope to be able to continue at the same rate. Within the constraints of space (the references should be consulted for precise mathematical statements and results), Section 2.2 will provide some more discussion of ISS and related concepts.

But the story is by no means complete. Among ongoing work, we may cite the following: (a) The formulation of yet more systems properties in ISS-like terms, including a very general notion of minimum-phase systems and applications to adaptive control (cf. [44] for preliminary work), and more generally the use of ISS, iISS, etc., ideas in the analysis of performance of switching controllers (work by Morse, Hespanha, and others). (b) The notion of IOSS was originally formulated in the context of understanding output stabilization problems, but the existence of observers hinges upon *incremental* IOSS (cf. [2]). This is very closely related to the OSS-type detectability notions pursued in work by Krener, and the PI's graduate student Ingalls (he will be completing his thesis next summer) has made major advances on this subject. (c) We have begun working-out of details, including Lyapunov converse theorems, for notions combining IOSS and IOS (a paper is in preparation), which reflect a "measurement to output" stability condition that appears naturally in regulation problems. (d) Applications of notions of IOS to the enhancement of nonlinear regulation theory so as to also include the analysis of unmodeled exosystem components (cf. [12] for preliminary thoughts).

1.2 Small Disturbances

Returning to the general theme of external signals d and their influence on system behavior, recall that we distinguished between "large d " gain analysis, which we carry out in an ISS-like setting, and "small d " (such as observation errors) which may lead to degradation of performance and chattering behavior. The latter phenomena, to which we turn next, have always been understood at an intuitive level, but we firmly believe that current research – which builds upon fundamental work carried out by Filippov, Krasovskii, Subbotin, Hermes, Hájek, and others in the 1960s and 1970s, as well as Brockett, Coron, Rosier, Ryan, and many others in the 1980s and early 1990s – has resulted in a deeper conceptual understanding, and certainly on very precise and more definitive mathematical formulations and theorems.

Many different issues could be discussed here, but, for concreteness (and because this is to where we have directed a major part of our past efforts) let us restrict ourselves to the problem of stabilizing control systems

$$\dot{x} = f(x, u)$$

under state feedback subject to vanishingly small measurement noise. We make standard technical hypotheses: states $x(t)$ evolve in \mathbb{R}^n (or, more generally, a manifold) controls take values $u(t)$ in a metric space \mathcal{U} , f is continuous and locally Lipschitz on x uniformly on compact subsets of $\mathbb{R}^n \times \mathcal{U}$. Given is a compact subset \mathcal{A} of \mathbb{R}^n to be stabilized, such as an equilibrium or a desired periodic orbit of the unforced system. We will focus on state feedback; of course, one might equally well study questions involving dynamic output feedback.

While most typically in control theory one formulates the stabilization problem as that of finding a feedback law $k : \mathbb{R}^n \rightarrow \mathcal{U}$ with the property that the set \mathcal{A} becomes globally asymptotically stable with respect to the closed-loop system $\dot{x}(t) = f(x(t), k(x(t)))$, in practice observation errors due to quantization or noise are often unavoidable. Under such circumstances, a more appropriate closed-loop model is obtained when the argument in the feedback law is a corrupted measurement, i.e. $\dot{x}(t) = f(x(t), k(x(t) + e(t)))$, where $e(t)$ is an unknown observation error. (In digital control systems, interactions between time *sampling* and quantization introduce severe additional complications. These interactions are ignored in the discussion to follow, since we do not yet have anything interesting to say in that regard. Note, however, that the PI wrote several papers on nonlinear sampled data systems in the mid 1980s, and recently coauthored [10]. In addition, the PI supervises students – a Master's thesis was recently completed – in experiments to elucidate the effect of bandwidth limitations on measurement hardware which relies upon quantizers. This experimental work is being carried out at the PI's Control Laboratory, and has the eventual goal of formulating meaningful interesting research problems.)

For systems subject to small observation error, a reasonable objective is to ask if trajectories converge to \mathcal{A} *provided that the errors $e(t)$ are sufficiently small with respect to distance to the target*. As a trivial illustration of what we mean, take the problem of stabilization of linear systems $\dot{x} = Ax + Bu$ to $\mathcal{A} = \{0\}$ via linear feedback $k(x) = -F(x + e)$. In this case, if $x'Px$ is a quadratic Lyapunov function, where $(A - BF)'P + P(A - BF) = -I$ and $P = P' > 0$, and writing $d := -BF e$, we have that $\dot{x} = (A - BF)x + d$, so $d(x'Px)/dt = -|x|^2 + 2x'Pd \leq -|x|^2/2$ whenever $|d| \leq |x|/(4\|P\|)$, and thus stability holds with a “margin of measurement error” $|e(t)| \leq |x(t)|/(4\|P\|\|BF\|)$. That is, the tolerated error (such that decrease of the Lyapunov function is insured) is small when near the set \mathcal{A} (the origin, in this case). *In general, a similar argument holds for arbitrary nonlinear systems*, if we are given a state-feedback law k that is *continuous*. A sketch of proof is as follows. Suppose that V is a Lyapunov function for verifying the asymptotic stability of \mathcal{A} with respect to the closed-loop system $\dot{x} = f(x, k(x))$; thus, $\nabla V(x) \cdot f(x, k(x)) < 0$ for all $x \notin \mathcal{A}$. Since

$$\nabla V(x) \cdot f(x, k(x + e)) \approx \nabla V(x) \cdot f(x, k(x)),$$

we have that also $\nabla V(x) \cdot f(x, k(x + e)) < 0$ for $x \in \mathbb{R}^n \setminus \mathcal{A}$, as long as e is small enough so that the approximation holds. This “small enough” condition can be expressed in the form “ $|e(t)| \leq \gamma(|x(t)|_{\mathcal{A}})$ for all t ” where $|\cdot|_{\mathcal{A}}$ denotes distance to \mathcal{A} , and where γ is some positive definite continuous function which can be interpreted as a robustness margin for the stabilization problem. Note that *this proof breaks down when k is not continuous*, because in that case the vector $f(x, k(x + e))$ may point in a very different direction than $f(x, k(x))$. There is however an alternative, which does not rely upon continuity of k but, instead, uses continuous differentiability of V . The approximation

$$\nabla V(x) \cdot f(x, k(x + e)) \approx \nabla V(x + e) \cdot f(x + e, k(x + e))$$

allows one to invoke the Lyapunov property *at the point $x + e$ instead of at x* ; the last expression will be negative as long as e is sufficiently small (for instance, when $\mathcal{A} = \{0\}$, if $x + e \neq 0$), and this leads to a robust stabilization theorem when there exists a \mathcal{C}^1 V . We are being deliberately vague, in this intuitive discussion, regarding the meaning of “solution” of the differential equation $\dot{x} = f(x, k(x + e))$, a nontrivial issue when k is not continuous. Precise formulations and proofs can be found in [5], and we also include some more details in Section 2.1.

To summarize, closed-loop convergence to \mathcal{A} is guaranteed, for small enough measurement errors, if the original feedback law (for the system not subject to measurement errors) is continuous, or if it is not necessarily continuous but stability can be verified by some continuously differentiable Lyapunov function. However, and these are key facts, *discontinuous feedback laws k are often unavoidable*, and for discontinuous k 's which stabilize, C^1 Lyapunov functions V do not always exist.

The study of *necessity* of discontinuous k should not be confused the fact that many optimal control problems result in discontinuous, e.g. “bang-bang” feedback. In contrast to optimal controls, the stabilization problem is highly underdetermined, a feedback law being basically an arbitrary section of a suitable map, so the question of existence of continuous stabilizing feedback is very geometric/topological in character. The C^0 feedback problem has been the subject of a substantial research effort during the last 20 years, including the counterexamples and comparisons between static and dynamic feedback in the PI’s paper with Sussmann in 1980, and the celebrated Brockett necessary condition as well as major later work by Coron and others. The obstructions to continuity of global stabilizers are obvious when the state space is not Euclidean and the target set \mathcal{A} is an equilibrium (Milnor’s theorem tells us that the domain of attraction of an equilibrium must be diffeomorphic to Euclidean space), or if the state space is \mathbb{R}^n but \mathcal{A} is not simply connected, such as a periodic orbit in \mathbb{R}^2 or even the set $\mathcal{A} = \{-1, 1\}$ in \mathbb{R} (for the latter, take the one-dimensional system $\dot{x} = u$ with $\mathcal{U} = \mathbb{R}$; this set can be made globally asymptotically stable under many different discontinuous feedback laws, such as $k(x) = 1 - x$ for $x \geq 0$ and $k(x) = -1 - x$ for $x < 0$, but there is no possible continuous global feedback stabilizer $k(x)$ because the domain of attraction of \mathcal{A} must be homotopically equivalent to \mathcal{A} , which has two connected components, and hence cannot be all of \mathbb{R}). But even if \mathcal{A} is merely an equilibrium point, let us say $\mathcal{A} = \{0\}$, the state space is Euclidean, and only local stabilization is desired, it may still be impossible to stabilize continuously, due to the existence of “virtual obstacles” imposed by restrictions upon the possible infinitesimal directions of motion (nonholonomy). This is the gist of Brockett’s analysis, applied e.g. to the usual car or unicycle models. Let us discuss an example which will lead us naturally back into the effect of measurement errors.

Consider the system $\dot{x} = f(x, u) = g(x)u = (x_1^2 - x_2^2, 2x_1x_2)'u$ with state space \mathbb{R}^2 and input space $\mathcal{U} = \mathbb{R}$ (if we think of points in \mathbb{R}^2 as complex numbers z , the system may be described simply as $\dot{z} = z^2u$). The orbits of the vector field g are the circles tangent to the x_1 axis and centered on the x_2 -axis, as well as $\{0\}$ and the positive and negative x_1 axes. The vector field g and typical orbits of g are shown in Figure 1. On the circles, input values $u(t)$ that are positive,




Figure 1: vector field $(x_1^2 - x_2^2) \frac{\partial}{\partial x_1} + 2x_1x_2 \frac{\partial}{\partial x_2}$ with typical integral manifolds

negative, or zero correspond respectively to motions that go counterclockwise or clockwise along

the integral curves of g , or stay fixed. Suppose that we wish to asymptotically stabilize the origin. As no motion can ever exit the integral curves, we will simply study what happens in a fixed circle, say one on the top half-plane. We pick a circle small enough to be in the domain of attraction of the closed-loop system. (We could think of the system restricted to our circle as a system on \mathbb{S}^1 given in polar coordinates by $\dot{\theta} = u$, that is, controls allow us to move at any speed, clockwise or counterclockwise. Then, Milnor's theorem already tells us that *continuous* stabilization is impossible, as \mathbb{S}^1 is not diffeomorphic to \mathbb{R} . But we wish to argue in much more elementary terms.) Let us call p the bottom point in the circle (the point $0 \in \mathbb{R}^2$). Near p , and to its right, stability says that $u = k(x)$ must make us move clockwise, that is, $k(x) < 0$, while near but to the left of p we must move counterclockwise, i.e. $k(x) > 0$. If k were to be a continuous function of the state x , the intermediate value theorem would force the existence, on this circle, of some state $q \neq p$ at which $k(q) = 0$. Thus q is another equilibrium point, and so k cannot asymptotically stabilize p .

So a continuous feedback stabilizer cannot exist, which means that the argument given earlier for insensitivity to measurement errors cannot be applied (also, one can prove that no C^1 Lyapunov function V exists). On the other hand, the system is asymptotically controllable (every state can be driven, in infinite time, to the origin), and the general theorem given in [1] says that feedback stabilizers exist for all such systems.* Indeed, there are for this example many obvious discontinuous feedback stabilizers. The most obvious solution is to go clockwise if $x_1 > 0$ and counterclockwise if $x_1 < 0$, making an arbitrary choice, let us say clockwise, if $x_1 = 0$. When restricting to some fixed circle as above, this would mean that one moves each state to the right of the top point p' , including p' itself, clockwise towards p , and every state to the left of p' counterclockwise towards p , cf. Figure 2(a). However, this feedback law is extremely




Figure 2: (a) feedback on circles; (b) true state, measured state, erroneous motion

sensitive to measurement errors, and may result in chattering behavior at the discontinuity p' . Indeed, if the true state x is slightly to the left of the top point p' , but our measurements make us believe it to be to the right, we will mistakenly use a clockwise motion, in effect bringing the state towards the top, instead of downwards towards the target (the origin); see Figure 2(b). It is clear that, if we are unlucky enough to consistently make measurement errors that place us on the opposite side of the true state, the result may well be an oscillation around the top. So the feedback law which we gave is not “robust” with respect to measurement errors, in that it may get “stuck” due to such errors. To emphasize this point, we show in Figure 3(a) the result of a straightforward numerical simulation. On the vertical axis of the graph, 0 corresponds to the top point in the circle, and 1, -1 are both identified with the bottom point (target). (We used random, instead of consistently worst-case, small observation noise, to make the behavior even more striking, but of course, with random noise eventually the state must converge to the set $\{1, -1\}$ with probability one, as a simple random walk argument shows. However, we later contrast this transient behavior with the one in (b) resulting from a different technique.)

Not only is the one given feedback stabilizer highly sensitive to measurement errors, but there is *no possible* static state feedback which stabilizes robustly, in the sense of tolerating

*The paper [1] constituted a major achievement of the previous grant, and generated considerable attention and follow-up work by many other researchers. It was subsequently selected by the American Mathematical Society for a “featured review” in Math Reviews, a distinction awarded to a small fraction of papers in mathematics.




Figure 3: random noise simulation (a) state feedback; (b) hybrid hold technique

“small” noise. As a matter of fact, one can establish a *general theorem providing necessary and sufficient conditions for static state feedback stabilization under small observation error for arbitrary systems* (and the theorem applies of course to the above relatively trivial example); this is another one of the current grant’s achievements, cf. [5]. For systems affine in controls, the basic principle (Hermes, Coron and Rosier, Ryan)) is that stability under small noise implies stability of all solutions to an associated differential inclusion (technically, Krasovskii or Filippov solutions, depending on whether one takes into account or not behavior on sets of measure zero), and solutions for any initial state constitute a closed connected set, leading to the same topological obstructions as those forced by continuous feedback; for general nonlinear systems, the proof relies upon showing that robust stabilization under small noise implies the existence of a smooth Lyapunov function. Section 2.1 has a brief discussion, and [5] provides details.

An Alternative: Hybrid Dynamic Feedback

Techniques which are commonly used in engineering in order to avoid chattering, such as the introduction of deadzones, may be interpreted as dynamic (i.e., memory-based) feedback. This observation, coupled with ideas originating in differential games and due to Subbotin and Krasovskii, led us into the general theorem presented in our 1997 CISS paper with Ledyayev: *given a stabilizing feedback, it is always possible to modify it in such a manner that closed-loop stabilization becomes robust to small observation errors*. Our modified controller incorporates an internal model of the system, and compares, at appropriate sampling times, the state predicted by the internal model with the noisy observations of the state; whenever these differ substantially, a “resetting” is performed on the state of the controller. More recently, in another noteworthy contribution done under current support, we came up with an alternative version of the same general idea, but one that is simpler and requires relatively little work beyond that (for stabilization in the absence of observation noise) found in [1]. The corresponding theorem is stated and proved in [13]; let us explain the basic principle now but defer more discussion to Section 2.1.

The basic idea of our approach to robustness under small observation errors is a kind of hysteresis-type behavior, namely to “stay the course” required by the controller for a certain minimal time, which depends on assumed upper bounds on the observation error, and to only re-estimate the state, and the resulting value of the feedback law, after this minimal time period has elapsed. Starting at some initial time, let us say $t_0 = 0$, we measure $k(x(0)+e(0))$, obtaining a control value u_0 . We now apply the constant control $u(t) \equiv u_0$ during a small interval $[0, t_1)$.

This procedure is then repeated: at time t_1 we resample, obtaining some $u_1 := k(x(t_1) + e(t_1))$, and apply this control during another small interval $[t_1, t_2)$, and so on. This modified strategy is not a pure continuous time one, in that a minimum intersample time is required, nor is it a standard sample-and-hold strategy, since the sampling rate is not necessarily constant. It is a “hybrid” control law, in the sense that both discrete and continuous time are involved, and logical decisions are taken.

The resulting behavior can be illustrated with the simple example of a motion in a circle discussed earlier. The controller now decides in which direction to move, clockwise or counterclockwise, based on the current measurement (which may well be somewhat inaccurate), but it forces the system to continue moving in the chosen direction, *not paying any attention to observations*, for a small time interval. Only after that interval does it measure again and recompute the control direction. Figure 3(b) shows the result of a simulation when this “don’t sample too fast” feature is added; note the quick convergence to the bottom point p (1 and -1 on the vertical axis). Of course, this example is much too simple to reflect the complexity of the general result, because there is only one discontinuity in the feedback (at p'). In general, the only available feedback law may well have an intricate set of discontinuities. Though simple to explain, proving that one can always do this successfully, for any asymptotically controllable system (i.e., for any system for which it is theoretically possible to control to the target) requires careful estimates of decrease for appropriate control-Lyapunov functions, cf. [13].

1.3 Control-Lyapunov Functions

Our constructions of feedback stabilizers, for general nonlinear systems in [1] as well as when adding observation errors, typically rely upon the use of *control-Lyapunov functions* (clf’s). Roughly speaking, a clf (relative to a target set \mathcal{A}) is a nonnegative function V of states such that $V(x) \approx 0$ when x is near \mathcal{A} and $V(x)$ is large when x is far from \mathcal{A} , with the property that $V(x(t))$ can be made to decrease, whenever away from \mathcal{A} , by an appropriate control action. If there is a clf, one may drive states towards the target by successive solution of a sequence of static nonlinear programming problems. (Depending on the application, one might dualize and attempt to maximize, rather than minimize, a measure of success.) At this vague level, clf’s include classical Lyapunov functions, value (Bellman) functions in optimal control, geodesic distance to targets in geometry, position evaluators in A.I. game-playing programs, critics (e.g. via neural-networks) in learning control, and even computer science objects such as Floyd/Dijkstra’s “variant” used to verify program termination.

A *precise* definition of clf’s for arbitrary finite-dimensional continuous-time systems was given in the PI’s 1982 CLF paper, expressing the decrease condition using Dini derivatives. Equivalent definitions can be formulated in terms of minimization of directional viscosity (or proximal) derivatives, or using integral inequalities (see for instance the exposition [24]). It was shown in the 1982 paper that for stability to equilibria (the same proof works for general compact attractors, see [11]), the *existence of a continuous clf is equivalent to asymptotic controllability (ac)* to the equilibrium in question. Asymptotic controllability means, roughly, that each state can be driven, by means of some open-loop control action, to the equilibrium, but the precise formulation, especially regarding stable behavior around the target, is more delicate and was also elucidated in the 1982 paper. (Actually, an equivalent but far more convenient definition can be given in terms of \mathcal{KL} functions, see [24].)

Independently of our work in 1982, in the landmark and simultaneous paper, Artstein showed that *existence of a smooth clf is equivalent, for systems affine on controls, to existence*

of a continuous feedback stabilizer. (For non-affine systems, he provided a feedback which takes values on a convexified velocity set. A different necessary and sufficient characterization for non-affine systems was recently provided in [5], where we showed that smooth clf's exist *if and only if there is a feedback which is insensitive to small measurement errors.*) This insight tied together the work on existence of continuous feedback with the search for clf's, and constitutes a cornerstone of modern nonlinear control design. In the late 1980s, the PI introduced a small enhancement to Artstein's theorem, observing that it is possible to give a universal formula for a stabilizing feedback expressed in terms of directional derivatives of a known clf.

2 Some More details on Selected Topics

There is, obviously, no space in this report to provide details on more than a few selected topics, which we do next. For other topics, we refer the reader to the following Web page:

<http://www.math.rutgers.edu/~sontag/>

which provides access to a substantial number of papers by the PI.

2.1 Stability and Small Disturbances

For concreteness, we consider here exclusively continuous-time systems evolving in finite-dimensional spaces \mathbb{R}^n , and we suppose that controls take values in $\mathcal{U} = \mathbb{R}^m$. A control or input is any measurable locally essentially bounded map $u(\cdot) : [0, \infty) \rightarrow \mathcal{U}$. In general, we use the notation $|x|$ for Euclidean norms, and use $\|u\|$, or $\|u\|_\infty$ for emphasis, to indicate the essential supremum of a function $u(\cdot)$. Given a map $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ which is locally Lipschitz and satisfies $f(0, 0) = 0$, we consider the system

$$\dot{x}(t) = f(x(t), u(t)) \quad (1)$$

and, when f does not depend on u (for instance, when we substitute, later, a feedback law $u = k(x)$, leading to $\dot{x} = f(x, k(x))$), we have a system with no inputs

$$\dot{x}(t) = f(x(t)). \quad (2)$$

All definitions for such systems are implicitly applied as well to systems with inputs (1) by setting $u \equiv 0$; for instance, we say that (the equilibrium solution $x \equiv 0$ of) (1) is globally asymptotically stable (GAS) if this holds for $\dot{x} = f(x, 0)$. The maximal solution $x(\cdot)$ of (1), corresponding to a given initial state $x(0) = x^0$, and to a given control u , is defined on some maximal interval $[0, t_{\max}(x^0, u))$ and is denoted by $x(t, x^0, u)$ ($x(t, x^0)$ for systems with no inputs (2)) or just $x(t)$ if clear from context.

The use of “comparison functions” has become widespread in stability analysis, as this formalism allows elegant formulations of most concepts. We recall the relevant definitions here. The class of \mathcal{K}_∞ functions consists of all $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ which are continuous, strictly increasing, unbounded, and satisfy $\alpha(0) = 0$. The class of \mathcal{KL} functions consists of those $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ with the properties that (1) $\beta(\cdot, t) \in \mathcal{K}_\infty$ for all t , and (2) $\beta(r, t)$ decreases to zero as $t \rightarrow \infty$. In [4], we proved that, for each $\beta \in \mathcal{KL}$, there exist two functions $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$ so that

$$\beta(r, t) \leq \alpha_2(\alpha_1(r)e^{-t}) \quad (3)$$

for all s, t . (Thus, as every function of the latter form is in \mathcal{KL} , and since \mathcal{KL} functions are only used in order to express upper bounds, in a sense there is no need to introduce the class \mathcal{KL} .) We will also use \mathcal{N} to denote the set of all nondecreasing functions $\sigma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$.

Expressed in this language, the property of *global asymptotic stability* (GAS) for a system with no inputs (2) becomes:

$$(\exists \beta \in \mathcal{KL}) \quad |x(t, x^0)| \leq \beta(|x^0|, t) \quad \forall x^0, \forall t \geq 0.$$

It is an easy exercise to show that this definition is equivalent to the usual “ $\varepsilon - \delta$ ” definition.

2.1.1 Stabilization

Recall that a system with inputs (1) is (open loop, globally) *asymptotically controllable* (AC) if for each initial state x^0 there exists some control $u = u_{x^0}(\cdot)$ defined on $[0, \infty)$, such that the corresponding solution $x(t, x^0, u)$ is defined for all $t \geq 0$, and converges to zero as $t \rightarrow \infty$ with “small” overshoot, and $u(t)$ does not become unbounded for x near zero. The precise formulation is as follows.

$$(\exists \beta \in \mathcal{KL}) (\exists \sigma \in \mathcal{N}) \quad \forall x^0 \in \mathbb{R}^n \quad \exists u(\cdot), \|u\|_{\infty} \leq \sigma(|x^0|),$$

$$|x(t, x^0, u)| \leq \beta(|x^0|, t) \quad \forall t \geq 0.$$

Informally, we say that $k : \mathbb{R}^n \rightarrow \mathcal{U}$ is a *feedback stabilizer* for the system with inputs (1) if k is locally bounded (that is, k is bounded on each bounded subset of \mathbb{R}), $k(0) = 0$, and the closed-loop system

$$\dot{x} = f(x, k(x)) \tag{4}$$

is GAS, i.e. there is some $\beta \in \mathcal{KL}$ so that $|x(t)| \leq \beta(|x(0)|, t)$ for all solutions and all $t \geq 0$. Obviously, if there exists a feedback stabilizer for (1), then (1) is also AC.

The reason that this is not yet a precise definition is that we have not defined, nor is there a standard notion, of “solution” of an initial value problem for (1) when $f(x, k(x))$ is not continuous. Various weak forms of solutions have been proposed, and we will discuss one of them soon. Leaving aside for now this (important) technical point, a most natural question is to ask if the converse holds as well, namely: *is every asymptotically controllable system also feedback stabilizable?* We will see that the answer is “yes”, provided that we allow discontinuous feedbacks k . In control problems, there are often ultimate objectives other than stabilization. For example, one might merely be interested in making a *subset* of the state variables (as opposed to the full state) approach a desired set point, in which case (after translation of the origin of coordinates so as to make the desired values zero) one deals with partial or “output” stability: there is a given function $y = h(x)$ and one wishes to make $y(t) \rightarrow 0$ using suitable control actions. A variation of this idea is the formulation of tracking problems, where one wishes the output $y(t)$ to follow a predetermined path $y^*(t)$; this can be seen as an output stability problem for an “error” signal $e(t) = y(t) - y^*(t)$. Yet another variation consists in studying stability, perhaps only local or with respect to the complement of a measure-zero set, of invariant sets (as opposed to equilibria); for instance, one may be given a periodic orbit, and a feedback is sought so that for the closed-loop system it becomes a limit cycle. Finally, one may wish to allow for uncertainties in plant description (adaptive and/or robust control). Many such other objectives can be treated using the techniques that we shall describe; we focus on stabilization to an equilibrium (or to an invariant compact set; except for notations, that

case if not much different) because this is the simplest problem in which the main technical difficulties can be understood.

We now return to the discussion in Section 1.2, and make precise some of the definitions and results cited there. We say that a continuous function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is positive definite if $V(x) = 0$ only if $x = 0$, and it is proper (or “weakly coercive”) if for each $a \geq 0$ the set $\{x \mid V(x) \leq a\}$ is compact, or, equivalently, $V(x) \rightarrow \infty$ as $|x| \rightarrow \infty$ (radial unboundedness). A property which is equivalent to properness and positive definiteness together is:

$$(\exists \underline{\alpha}, \bar{\alpha} \in \mathcal{K}_\infty) \quad \underline{\alpha}(|x|) \leq V(x) \leq \bar{\alpha}(|x|) \quad \forall x \in \mathbb{R}^n. \quad (5)$$

A continuous $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is a *control-Lyapunov function* (clf) if it is proper, positive definite, and infinitesimally decreasing in the following generalized sense: there exist a positive definite continuous $W : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ and a $\sigma \in \mathcal{N}$ so that

$$\sup_{x \in \mathbb{R}^n} \max_{\zeta \in \partial_P V(x)} \min_{|u| \leq \sigma(|x|)} \zeta \cdot f(x, u) + W(x) \leq 0 \quad (6)$$

where $\partial_P V(x)$ is the set of all proximal subgradients at x . An equivalent property is to ask that V be a *viscosity supersolution* of the corresponding Hamilton-Jacobi-Bellman equation. In the paper the PI’s 1982 CLF paper, the PI showed that *a system (1) is asymptotically controllable if and only if it admits a (continuous) clf*. (Not surprisingly, the proof is based on first constructing an appropriate W , and then letting V be the optimal cost (Bellman function) for the problem $\min \int_0^\infty W(x(s)) ds$; however, some care has to be taken to insure that V is continuous, and the cost has to be adjusted in order to deal with possibly unbounded minimizers. To be precise, the result as stated here is really a restatement of the main theorem given in the PI’s 1982 CLF paper, cf. [24] for a discussion.)

In the more standard continuously differentiable case, the definition reduces to the usual one that for each nonzero $x \in \mathbb{R}^n$ there is some u with $|u| \leq \sigma(|x|)$ such that $\nabla V(x) \cdot f(x, u) < 0$. As discussed in Section 1.3, Artstein showed that existence of a smooth clf is equivalent, for systems affine on controls, to existence of a *continuous* feedback stabilizer. The original Arstein’s proof proceeds by a nonconstructive argument involving partitions of unity, but it is also possible to exhibit explicitly a feedback, written as a function $k(\nabla V(x) \cdot g_0(x), \dots, \nabla V(x) \cdot g_m(x))$ of the directional derivatives of V along the vector fields defining the system (*universal formulas* for stabilization). Taking for simplicity $m = 1$, one such formula is:

$$k(x) := - \frac{a(x) + \sqrt{a(x)^2 + b(x)^4}}{b(x)} \quad (0 \text{ if } b = 0)$$

where $a(x) := \nabla V(x) \cdot g_0(x)$ and $b(x) := \nabla V(x) \cdot g_1(x)$. (The expression for k is analytic in a, b when $x \neq 0$, because the clf property means that $a(x) < 0$ whenever $b(x) = 0$. Continuity at zero is achieved by the formula if a “small control property” holds.)

However, it has been well-known, at least since the work by Sussmann, that it is impossible to stabilize an arbitrary AC system using continuous feedback. Brockett’s condition provides a fundamental necessary condition for existence of continuous stabilizers; the basic fact, due to Krasnosel’ski, is that if the system $\dot{x} = F(x) = f(x, k(x))$ has the origin as an asymptotically stable point then the degree (index) of F with respect to zero is $(-1)^n$, where n is the system dimension. See e.g. [24]. (We showed in work with Sussmann in the early 1980s that every one-dimensional AC system can be stabilized using a continuous time-varying feedback $u = k(t, x)$, and it is not difficult to see that the proof can be adapted to obtain $k(t, x)$ periodic in t , for

each x . Subsequently, Coron established a general result on periodic time-varying stabilization valid in any dimension, but restricted to systems with *no drift*. He later extended the results to local stabilization of systems with drift which satisfy suitable Lie-algebraic local controllability conditions, and with Rosier also showed that such feedbacks exist provided that there are smooth clf's. But an example of Sussmann's shows that time varying, or more general dynamic, continuous feedback stabilizers do not exist in general.)

Thus, in order to develop a satisfactory general theory of stabilization, one in which one proves the implication "asymptotic controllability implies feedback stabilizability," one must allow discontinuous feedback laws $u = k(x)$. But then, a major technical difficulty arises: solutions of the initial-value problem $\dot{x} = f(x, k(x))$, $x(0) = x^0$, interpreted in the classical sense of differentiable functions or even as (absolutely) continuous solutions of the associated integral equation do not exist in general. There is, of course, an extensive literature addressing the question of discontinuous feedback laws for control systems and, more generally, differential equations with discontinuous right-hand sides. One of the best-known candidates for the concept of solution of (4) is that of a *Filippov solution*, which is defined as the solution of a certain differential inclusion with a multivalued right-hand side which is built from $f(x, k(x))$. Unfortunately, there is no hope of obtaining the implication "asymptotic controllability implies feedback stabilizability" if one interprets solutions of (4) as Filippov solutions. This is a consequence of results by Ryan, Coron, and Rosier, which established that the existence of a discontinuous stabilizing feedback in the Filippov sense implies the Brockett necessary conditions, and, moreover, for systems affine in controls it also implies the existence of continuous feedback (which we know is in general impossible).

A different concept of solution originates with the theory of discontinuous positional control developed by Krasovskii and Subbotin in the context of differential games, and it is the basis of the new approach to discontinuous stabilization proposed in [1], based in interpretation of feedback as *limits of high-frequency sampling*. A *sampling schedule* $\pi = \{t_i\}_{i \geq 0}$ is an infinite sequence $0 = t_0 < t_1 < t_2 < \dots$ with $\lim_{i \rightarrow \infty} t_i = \infty$. We call $\bar{d}(\pi) := \sup_{i \geq 0} (t_{i+1} - t_i)$ the *diameter* of π . Suppose that k is a given feedback law for system (1). For each π , the π -*trajectory starting from* x^0 of system (4) is defined recursively on the intervals $[t_i, t_{i+1})$, $i = 0, 1, \dots$, as follows. On each interval $[t_i, t_{i+1})$, the initial state is measured, the control value $u_i = k(x(t_i))$ is computed, and the constant control $u \equiv u_i$ is applied until time t_{i+1} ; the process is then iterated. That is, we start with $x(t_0) = x^0$ and solve recursively

$$\dot{x}(t) = f(x(t), k(x(t_i))), \quad t \in [t_i, t_{i+1}), \quad i = 0, 1, 2, \dots$$

using as initial value $x(t_i)$ the endpoint of the solution on the preceding interval (this is *not* the same as an Euler-polygon solution, since $x(t)$, not $x(t_i)$, appears in the right-hand side in the first argument). The ensuing π -trajectory, which we denote as $x_\pi(\cdot, x^0)$, is defined on some maximal nontrivial interval; it may fail to exist on the entire interval $[0, +\infty)$ due to a blow-up on one of the subintervals $[t_i, t_{i+1})$. We say that it is *well defined* if $x_\pi(t, x^0)$ is defined on all of $[0, +\infty)$.

Definition. The feedback $k : \mathbb{R}^n \rightarrow \mathcal{U}$ *stabilizes* the system (1) if there exists a function $\beta \in \mathcal{KL}$ so that the following property holds: For each $0 < \varepsilon < K$ there exists a $\delta = \delta(\varepsilon, K) > 0$ such that, for every sampling schedule π with $\bar{d}(\pi) < \delta$, and for each initial state x^0 with $|x^0| \leq K$, the corresponding π -trajectory of (4) is well-defined and satisfies

$$|x_\pi(t, x^0)| \leq \max \{ \beta(K, t), \varepsilon \} \quad \forall t \geq 0. \quad (7)$$

Observe that, in particular, one then has $|x_\pi(t, x^0)| \leq \max\{\beta(|x^0|, t), \varepsilon\}$ for all $t \geq 0$ whenever $0 < \varepsilon < |x^0|$ and $\bar{d}(\pi) < \delta(\varepsilon, |x^0|)$ (just take $K := |x^0|$). The role of δ is to specify a lower bound on intersampling times; one is requiring that $t_{i+1} \leq t_i + \theta(|x(t_i)|)$ for each i , where θ is an appropriate positive function.

Our definition of stabilization is physically meaningful, and is very natural in the context of sampled-data (computer control) systems. It says in essence that a feedback k stabilizes the system if it drives all states asymptotically to the origin and with small overshoot when using *any fast enough sampling schedule*. A high enough sampling frequency is generally required when close to the origin, in order to guarantee small displacements, and also at infinity, so as to preclude large excursions or even blow-ups in finite time. This is the reason for making δ depend on ε and K .

This concept of stabilization can be reinterpreted in various ways. One is as follows. Pick any initial state x^0 , and consider any sequence of sampling schedules π_ℓ whose diameters $\bar{d}(\pi_\ell)$ converge to zero as $\ell \rightarrow \infty$ (for instance, constant sampling rates with $t_i = i/\ell$, $i = 0, 1, 2, \dots$). Note that the functions $x_\ell := x_{\pi_\ell}(\cdot, x^0)$ remain in a bounded set, namely the ball of radius $\beta(|x^0|, 0)$ (at least for ℓ large enough, for instance, any ℓ so that $\bar{d}(\pi_\ell) < \delta(|x^0|/2, |x^0|)$). Because $f(x, k(x))$ is bounded on this ball, these functions are equicontinuous, so by Arzela-Ascoli we may take a subsequence, which we denote again as $\{x_\ell\}$, so that $x_\ell \rightarrow x$ as $\ell \rightarrow \infty$ (uniformly on compact time intervals) for some absolutely continuous (even Lipschitz) function $x : [0, \infty) \rightarrow \mathbb{R}^n$. We may think of any limit function $x(\cdot)$ that arises in this fashion as a *generalized solution of the closed-loop equation (4)*. That is, generalized solutions are the limits of trajectories arising from arbitrarily high-frequency sampling when using the feedback law $u = k(x)$. Generalized solutions, for a given initial state x^0 , may not be unique – just as may happen with continuous but non-Lipschitz feedback – but there is always existence; and, moreover, for any generalized solution, $|x(t)| \leq \beta(|x^0|, t)$ for all $t \geq 0$. This is precisely the defining estimate for the GAS property. (This type of interpretation is somewhat analogous, at least in spirit, to the way in which “relaxed” controls are interpreted in optimal trajectory calculations, namely through high-frequency switching of approximating regular controls.)

The following fundamental result was established in [1]: *The system (1) admits a stabilizing feedback if and only if it is asymptotically controllable*. Necessity is clear. The sufficiency statement is proved by construction of k , and is based on the following ingredients: (1) existence of a continuous clf V (already discussed); (2) regularization on shells of V ; and (3) pointwise minimization of a Hamiltonian for the regularized V . Once V is known to exist, the next step in the construction of a stabilizing feedback is to obtain Lipschitz approximations of V . For this purpose, one considers the Iosida-Moreau inf-convolution of V with a quadratic function:

$$V_\alpha(x) := \inf_{y \in \mathbb{R}^n} \left[V(y) + \frac{1}{2\alpha^2} |y - x|^2 \right]$$

where the number $\alpha > 0$ is a parameter which is constant on elements of a shell decomposition of the state-space. One has that $V_\alpha(x) \nearrow V(x)$, uniformly on compacts. The critical fact is that V_α is itself a clf for the original system when restricted to an appropriate shells. Since V_α is locally Lipschitz, it is differentiable almost everywhere (Rademacher), and roughly, the feedback k is then made equal to a pointwise minimizer k_α of the Hamiltonian, at the points of differentiability via $k_\alpha(x) := \text{any } u \in \mathcal{U}_0 \text{ minimizing } \nabla V_\alpha(x) \cdot f(x, u)$, where α and the compact $\mathcal{U}_0 = \mathcal{U}_0(\alpha)$ are chosen constant on shells. Actually, this description is oversimplified, and the proof is a bit more delicate. One must take the minimizers of $\zeta_\alpha(x) \cdot f(x, u)$, where $\zeta_\alpha(x)$ is carefully chosen. At points x of nondifferentiability, $\zeta_\alpha(x)$ is not a proximal subgradient of V_α ,

since $\partial_P V_\alpha(x)$ may well be empty. One uses, instead, the fact that $\zeta_\alpha(x)$ happens to be in $\partial_P V(x')$ for some $x' \approx x$.

We now return once more to the discussion in Section 1.2, regarding observation errors. It turns out that one may use feedback constructed in [1] in such a way as to obtain behavior robust with respect to small observation errors, using the idea, illustrated with the circle, of not sampling again for some minimal period. This discussion may be formalized in several ways. We limit ourselves here to a theorem assuring semiglobal practical stability (i.e., driving all states in a given compact set of initial conditions into a specified neighborhood of zero). For any sampling schedule π , we denote $\underline{d}(\pi) := \inf_{i \geq 0} (t_{i+1} - t_i)$. If $e : [0, \infty) \rightarrow \mathbb{R}^n$ is any function ($e(t)$ is to be thought of as the state estimation error at time t), k is a feedback law, $x^0 \in \mathbb{R}^n$, and π is a sampling schedule, we define the solution of $\dot{x} = f(x, k(x + e))$, $x(0) = x^0$ as earlier, namely, recursively solving $\dot{x}(t) = f(x(t), k(x(t_i) + e(t_i)))$ with initial condition $x(t_i)$ on the intervals $[t_i, t_{i+1}]$. The feedback stabilizer that was constructed in [1] was defined by patching together feedback laws, denoted there as k_ν (where $\nu = (\alpha, r, R)$ is a triple of positive numbers, with $r < R$). A theorem is given in [13], showing that, for any asymptotically controllable system (1), there exists a function $\Gamma \in \mathcal{K}_\infty$ with the following property: For each $0 < \varepsilon < K$, there is a feedback of the type k_ν , and there exist positive $\delta = \delta(\varepsilon, K)$, $\kappa = \kappa(\varepsilon, K)$, and $T = T(\varepsilon, K)$, such that, for each sampling schedule π with $\bar{d}(\pi) \leq \delta$, each $e : [0, \infty) \rightarrow \mathbb{R}^n$ so that $|e(t)| \leq \kappa \underline{d}(\pi)$ for all $t \geq 0$, and each x^0 with $|x^0| \leq K$, the solution of the noisy system satisfies $|x(t)| \leq \Gamma(K)$ for all $t \geq 0$ and $|x(t)| \leq \varepsilon$ for all $t \geq T$.

This theorem insures that stabilization is possible if we sample “just right” (not too slow, so as to preserve stability, but also not too fast, so that observation errors do not cause chattering). It leaves open the theoretical question of precisely under what conditions is it possible to find a state feedback law which is robust with respect to small observation errors and which, on the other hand, is a continuous-time feedback, in the sense of arbitrarily fast sampling. We discussed in Section 1.2 the fact that this objective cannot always be met (e.g., for the circle problem), but that the existence of a \mathcal{C}^1 clf might be sufficient for guaranteeing that it can. Indeed, this is what happens, as was proved in the recent paper [5]. The main result from that paper can be summarized by saying that there is a feedback which stabilizes under small errors (in a manner made precise in the paper, and using the high-frequency sampling notion of generalized solutions) if and only if there is a \mathcal{C}^1 clf for the unperturbed system (1). (A somewhat analogous result holds for classical solutions, as shown by Hermes and Hajek.)

It is interesting to note that, as a corollary of Artstein’s Theorem, for control-affine systems $\dot{x} = g_0(x) + \sum u_i g_i(x)$ we may conclude that if there is a discontinuous feedback stabilizer that is robust with respect to small noise, then there is also one that is smooth on $\mathbb{R}^n \setminus \{0\}$. For non control-affine systems, however, there may exist a discontinuous feedback stabilizer that is robust with respect to small noise, yet there is no regular feedback. For example, consider the following system with $n = 3$ and $m = 1$: $\dot{x}_1 = u_2 u_3$, $\dot{x}_2 = u_1 u_3$, $\dot{x}_3 = u_1 u_2$. Here, there is a \mathcal{C}^1 clf, namely the squared norm $(x_1^2 + x_2^2 + x_3^2)$, but there is no possible regular feedback stabilizer, since Brockett’s condition fails because points $(0, \neq 0, \neq 0)$ cannot be in the image of $(x, u) \mapsto f(x, u)$. (Because this is a homogeneous system with no drift, Brockett’s condition rules out even feedbacks that are not continuous at the origin, see [5] for a remark to that effect.) An interesting research area is that of using the feedback construction from this paper to obtain a “backstepping” result for an example as the one just given (we actually arrived at this example while analyzing a model of underwater vehicles – cf. [43] – where a reduced system is driven by inputs u corresponding to the effect of linear on angular velocities, but we never returned to the original motivation).

The sufficiency part of the theorem from [5] proceeds by taking a pointwise minimization of the Hamiltonian. The necessity part is based on the following technical fact: if the perturbed system can be stabilized, then the differential inclusion $\dot{x} \in F(x) := \bigcap_{\varepsilon > 0} \overline{\text{co}} f(x, k(x + \varepsilon B))$ (where B denotes the unit ball in \mathbb{R}^n) is strongly asymptotically stable. One may then apply the recent converse Lyapunov theorem of Clarke, Ledyaev, and Stern for upper semicontinuous compact convex differential inclusions (these generalized to more general differential inclusions the theorem from our converse Lyapunov paper with Lin and Wang characterizing uniform asymptotic stability of systems with disturbances $\dot{x} = f(x, d)$) to deduce the existence of V .

The following table:

$$\begin{array}{ccccc} \mathcal{C}^1 V & \iff & \exists \text{ robust } k & & \\ \Downarrow & & \Downarrow & & \\ \mathcal{C}^0 V & \iff & \exists k & \iff & \text{AC} \end{array}$$

summarizes which implications hold; we write “robust” to mean stabilization of the system subject to observation and actuator noise.

2.2 Input to State Stability and Related Notions

In Section 1.1, we surveyed results and references, and discussed intuitively notions associated to input-to-state stability (ISS). Although by now well-covered in several textbooks, let us first review the basic concept of ISS with some more precision, so that we may explain some of the new directions which we have been pursuing more recently.

One of the original reasons for introducing the concept of ISS was the following often-remarked fact: even if a system (1) is GAS (recall that we mean by this that the trivial solution $x \equiv 0$ of (2) is GAS), there is no reason that, for solutions of (1) associated to *nonzero* controls, it should hold $u(t) \rightarrow 0$ implies $x(t) \rightarrow 0$ as $t \rightarrow \infty$ (CICS = converging-input converging-state property) or that $u(\cdot)$ bounded should imply that $x(\cdot)$ is necessarily bounded (BIBS = bounded-input bounded-state property).

These implications fail in general for nonlinear systems. For example, the solution of the one-dimensional single-input system $\dot{x} = -x + (x^2 + 1)u$, with initial state $x_0 = \sqrt{2}$, and input $u(t) = (2t + 2)^{-1/2} \rightarrow 0$, is the unbounded trajectory $x(t) = (2t + 2)^{1/2}$, in spite of the fact that the unforced system is GAS, and even worse, the bounded input $u \equiv 1$ results in a finite-time explosion. On the other hand, for linear (finite-dimensional) systems $\dot{x} = Ax + Bu$, these implications are always true. Not only that, but the explicit decay/gain estimate $|x(t)| \leq \beta(t) |x^0| + \gamma \|u_t\|_\infty$ holds, where $\beta(t) = \|e^{tA}\| \rightarrow 0$ and $\gamma = \|B\| \int_0^\infty \|e^{sA}\| ds$ for any Hurwitz matrix A , and u_t is the restriction of u to $[0, t]$ thought of as a measurable bounded function which is zero for $s > t$. The CICS and BIBS properties are immediate consequences of this explicit estimate.

In typical operator-theoretic stability analysis, one studies stability of input/state (or input/output) operators with respect to various norms, i.e., one looks for decay/gain estimates of the form shown above for linear systems. Notions like ISS and iISS arise naturally when one attempts to generalize linear gain analysis to nonlinear systems. There are two main norms which one usually imposes on input functions and state trajectories, namely square norms (to quantify energy) or supremum norms (to study worst-case behavior). This gives rise to the following three possible estimates (the case “ $L^\infty \rightarrow L^2$ ” is less interesting, being too restrictive):

$$[L^\infty \rightarrow L^\infty] \quad |x(t)| \leq c |x(0)| e^{-\lambda t} + c \sup_{s \in [0, t]} |u(s)| \quad \text{for all } t \geq 0, \quad (8)$$

$$[L^2 \rightarrow L^\infty] \quad |x(t)| \leq c|x(0)|e^{-\lambda t} + c \int_0^t |u(s)|^2 ds \quad \text{for all } t \geq 0, \quad (9)$$

$$[L^2 \rightarrow L^2] \quad \int_0^t |x(s)|^2 ds \leq c|x(0)| + c \int_0^t |u(s)|^2 ds \quad \text{for all } t \geq 0. \quad (10)$$

To be precise: there exist positive constants c and λ so that, for each measurable essentially bounded control u and each initial condition, the solution is defined for all t and it satisfies the indicated estimates. (These three properties are equivalent for *linear* systems, and in fact they amount to simply asking that the matrix A be Hurwitz.) We argued in Section 1.1 that notions of stability should be invariant under changes of variables. Let us see where this requirement leads us to.

By a *change of variables* in \mathbb{R}^q we mean a transformation $z = T(x)$, where $T : \mathbb{R}^q \rightarrow \mathbb{R}^q$ is a homeomorphism, $T(0) = 0$, and $T|_{\mathbb{R}^q \setminus \{0\}}$ is a diffeomorphism. Suppose that we take a state change of coordinates $x = T(z)$ and a change of input variables $u = S(v)$. Then, there are two functions $\underline{\alpha}, \bar{\alpha} \in \mathcal{K}_\infty$ so that $\underline{\alpha}(|z|) \leq |T(z)| \leq \bar{\alpha}(|z|)$ for all $z \in \mathbb{R}^n$, and, similarly, we can write $|S(v)| \leq \bar{\gamma}(|v|)$ for each $v \in \mathbb{R}^m$, for some $\bar{\gamma} \in \mathcal{K}_\infty$. For any input u and initial state x^0 , and the corresponding trajectory $x(t) = x(t, x^0, u)$, let us write $x(t) = T(z(t))$ and $u(t) = S(v(t))$ for all t , and $z^0 = z(0) = T^{-1}(x^0)$. Let us rewrite each of the estimates in terms of the z and v variables. Estimate (8) translates into: $\underline{\alpha}(|z(t)|) \leq ce^{-\lambda t}\bar{\alpha}(|z^0|) + c \sup_{s \in [0, t]} \bar{\gamma}(|v(s)|)$ for all $t \geq 0$. If we use again “ x ” and “ u ” for state and input variables, and we let $\beta(s, t) := \underline{\alpha}^{-1}(2ce^{-\lambda t}\bar{\alpha}(s))$ and $\gamma(s) := \underline{\alpha}^{-1}(2c\bar{\gamma}(s))$, we are led to an estimate as follows, where $\beta \in \mathcal{KL}$, $\gamma \in \mathcal{K}_\infty$:

$$|x(t)| \leq \beta(|x(0)|, t) + \sup_{s \in [0, t]} \gamma(|u(s)|). \quad (11)$$

(This would appear to be less general, because β has a very special form, but surprisingly, any \mathcal{KL} function β arises in this manner, cf. (3).) On the other hand, estimate (9) translates into: $\underline{\alpha}(|z(t)|) \leq ce^{-\lambda t}\bar{\alpha}(|z^0|) + c \int_0^t \bar{\gamma}(|v(s)|)^2 ds$ for all $t \geq 0$. With $\beta(s, t) := ce^{-\lambda t}\bar{\alpha}(s)$ and $\gamma(s) := c\bar{\gamma}(s)^2$, we have now an estimate as follows:

$$\underline{\alpha}(|x(t)|) \leq \beta(|x(0)|, t) + \int_0^t \gamma(|u(s)|) ds. \quad (12)$$

Finally, estimate (10) translates into: $\int_0^t \underline{\alpha}(|z(s)|) ds \leq c\bar{\alpha}(|z^0|) + c \int_0^t \bar{\gamma}(|v(s)|)^2 ds$ and with $\gamma(s) := c\bar{\gamma}(s)^2$ and $\kappa(s) := c\bar{\alpha}(s)$, invariance under change of variables in (10) leads to an estimate of the form:

$$\int_0^t \underline{\alpha}(|x(s)|) ds \leq \kappa(|x^0|) + \int_0^t \gamma(|u(s)|) ds. \quad (13)$$

We now define any system (1) to be *input to state stable (ISS)* if there exist functions $\beta \in \mathcal{KL}$ and $\underline{\alpha}, \gamma \in \mathcal{K}_\infty$ so that, for each initial state x^0 and input u (measurable and locally essentially bounded), the solution $x(t) = x(t, x^0, u)$ is defined for all $t \geq 0$, and the estimate (11) holds. Similarly, we say that a system (1) is *integral input to state stable (iISS)* provided that there exist functions $\beta \in \mathcal{KL}$ and $\underline{\alpha}, \gamma \in \mathcal{K}_\infty$ so that, for each initial state x^0 and input u (measurable and locally essentially bounded), the solution $x(t) = x(t, x^0, u)$ is defined for all $t \geq 0$, and the estimate (12) holds. Regarding the last estimate, that in Equation (13) (what one might call a “nonlinear H_∞ estimate” in analogy to the linear case), we do not give it a name, because, as remarked below, it is the same as ISS.

In summary, we have seen that the notions of ISS and iISS arise most naturally when one starts from standard stability notions for linear systems and applies nonlinear changes

of variables. (Analogously, although this is not usually mentioned, Lyapunov's asymptotic stability can be seen as the notion which arises, via change of state variables, from exponential stability.)

The estimate (11) implies that an ISS system must have a well-defined asymptotic gain, that is, there is some $\gamma \in \mathcal{K}_\infty$ so that $\limsup_{t \rightarrow +\infty} |x(t, x^0, u)| \leq \gamma(\|u\|_\infty)$ for all x^0 and u and also (take $u \equiv 0$) an ISS system has the property that the zero-input system (2) is GAS. The second property is independent of inputs, and the first involves no estimates on initial states. It is then surprising that ISS is in fact equivalent to the conjunction of these two properties, which represents a "superposition principle" for nonlinear control. Moreover, one can weaken the GAS assumption to simply (neutral) stability of the zero-input system (2), because the asymptotic gain property, when applied to the special case $u \equiv 0$, already implies that all trajectories of (2) converge to zero (attractivity property). The result, which is nontrivial, is as follows: *The system (1) is ISS if and only if it admits an asymptotic gain and the unforced system (2) is stable.* The basic difficulty is in establishing uniform convergence estimates for the states (i.e., constructing the β function). If the system were affine in u , the proof would follow from a more or less standard weak compactness argument; however, for general systems, the proof is far more delicate.

Another characterization of ISS is via Lyapunov-like notions. Let us say that a smooth $V : \mathbb{R}^n \rightarrow \mathbb{R}$ is an *ISS-Lyapunov function* for the system (1) if V is proper and positive definite, and, for some $\gamma, \alpha \in \mathcal{K}_\infty$,

$$\dot{V}(x, u) := \nabla V(x) \cdot f(x, u) \leq -\alpha(|x|) + \gamma(|u|) \quad \forall x, u. \quad (14)$$

That is, the dissipation inequality $V(x(t_2)) - V(x(t_1)) \leq \int_{t_1}^{t_2} w(u(s), x(s)) ds$ holds along all trajectories of the system, where, using Willems' terminology, the "supply" function is $w(u, x) = \gamma(|u|) - \alpha(|x|)$. The following fundamental result is from our earlier work with Wang: *The system (1) is ISS if and only if it admits an ISS-Lyapunov function.* The proof of sufficiency is by means of a simple comparison principle, but the other implication takes a bit more effort. The critical observation is that ISS implies "small gain robustness": if $u(t)$ is for all t sufficiently small compared to $x(t)$, then the term involving u can be dropped from the right-hand side of the ISS estimate. This leads to the property that there is some $\rho \in \mathcal{K}_\infty$ so that, defining $g(x, d) := f(x, d\rho(x))$, the system $\dot{x} = g(x, d)$ is asymptotically stable uniformly with respect to all measurable $d(\cdot) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$ with values in a ball of radius one. One may then invoke the converse Lyapunov theorem for robust stability proved in our work with Lin and Wang to obtain a smooth V so that $\nabla V(x) \cdot g(x, d) \leq -\alpha(x)$ for all x and d , and an ISS-Lyapunov function can be obtained by slightly modifying this V .


Recall that the " $L^2 \rightarrow L^2$ " stability property in (10) led to an estimate of the form (13), more precisely, for each x^0 and u the solution is well-defined for all $t \geq 0$ and (13) holds. It is shown in [4] that this property holds *if and only if the system is ISS*.

The " $L^2 \rightarrow L^\infty$ " linear stability property (9) led us to the iISS estimate (12). As with the ISS property, there is an elegant dissipation characterization. Let us call a smooth $V : \mathbb{R}^n \rightarrow \mathbb{R}$ an *iISS-Lyapunov function* for the system (1) if V is proper and positive definite, and, for some $\gamma \in \mathcal{K}_\infty$ and some *positive definite* α , $\dot{V}(x, u) \leq -\alpha(|x|) + \gamma(|u|)$ for all x, u (the difference with (14) is that here we do not require α to be class \mathcal{K}_∞). A recent result from [9] is as follows: *The system (1) is iISS if and only if it admits an iISS-Lyapunov function.* Notice that this implies (since a \mathcal{K}_∞ function is in particular positive definite) that every ISS system is also iISS, but the converse is false (for bounded but large u 's one may well have $\dot{V} > 0$). Recent research has resulted in other very illuminating characterizations. We call a system (1) together with

a continuous output function $y = h(x)$ (with $h(0) = 0$) *h-dissipative* if there exists a smooth, positive definite, and proper function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ such that $\dot{V}(x, u) \leq -\alpha(h(x)) + \sigma(|u|)$ for all x, u , and *weakly h-detectable* if, for all trajectories, $y(t) = h(x(t)) \equiv 0$ implies that $x(t) \rightarrow 0$ as $t \rightarrow \infty$. This is proved in [9]: *The system (1) is iISS if and only if there is some output h so that the system is weakly h-detectable and h-dissipative.*

Let us illustrate the iISS results through an application which, as a matter of fact, was the one that originally motivated much of the work in [9]. In a recent paper, Marino and Tomei proposed the reformulation of tracking problems by means of the notion of input to state stability. Their goal was to strengthen the robustness properties of tracking designs, and they found the notion of ISS to be instrumental in the precise characterization of performance. In fact, they emphasized the novelty of using the ISS notion in this role. It turns out, however, that a typical passivity-based tracking design may well *not* result in ISS behavior, as we illustrate now by means of an example in robotic control.

Consider the manipulator shown in Fig. 4(a). A simple model is obtained considering the



arm as a segment with mass M and length L , and the hand as a point with mass m . If we denote with r the position of the hand and with θ the angle of the arm, the equations for such a system are: $(mr^2 + ML^2/3)\ddot{\theta} + 2mr\dot{r}\dot{\theta} = \tau$, $m\ddot{r} - mr\dot{\theta}^2 = F$, where F and τ indicate external torques. We now study the closed-loop system which is obtained by choosing τ and F as:

$$\begin{aligned}\tau &= -k_{d1}\dot{\theta} - k_{p1}(\theta - \theta_d) \\ F &= -k_{d2}\dot{r} - k_{p2}(r - r_d),\end{aligned}\tag{15}$$

with $k_{p1}, k_{p2}, k_{d1}, k_{d2} > 0$. (For notational simplicity, we will also write $q = [\theta, r]^T$). This represents a typical passivity-based tracking design, when we think of r_d and θ_d as signals to be followed by r and θ .

Normally, one establishes tracking behavior, as well as the closed-loop stability of the system when the reference signal $q_d = (\theta_d, r_d)$ is constant; for such signals, one obtains $\dot{q} \rightarrow 0$ and $q \rightarrow q_d$ as $t \rightarrow +\infty$. In the spirit of input-to-state stability, however, it is natural to ask what is the sensitivity of the design to *additive measurement noise*. That is, suppose that the input applied to the system is, instead of (15): $\tau = -k_{d1}(\dot{\theta} + d_{11}) - k_{p1}(\theta + d_{12} - \theta_d)$, $F = -k_{d2}(\dot{r} + d_{21}) - k_{p2}(r + d_{22} - r_d)$, where the $d_{ij}(t)$'s are observation errors. The closed-loop system that results is then as follows:

$$\begin{aligned}(mr^2 + ML^2/3)\ddot{\theta} + 2mr\dot{r}\dot{\theta} &= u_1 - k_{d1}\dot{\theta} - k_{p1}\theta \\ u_2m\ddot{r} - mr\dot{\theta}^2 &= u_2 - k_{d2}\dot{r} - k_{p2}r\end{aligned}\tag{16}$$

with inputs $u_1 = -k_{d1}d_{11} - k_{p1}(d_{12} - \theta_d)$ and $u_2 = -k_{d2}d_{21} - k_{p2}(d_{22} - r_d)$, which we think of in its equivalent form as a four-dimensional system on states q, z , with $z = \dot{q}$. The goal stated earlier

is to qualitatively analyze the sensitivity of the full state (q, z) as a function of the measurement errors d_{ij} . As these errors are potentially arbitrary functions, this problem amounts to the study of stability properties with respect to arbitrary input functions $u = (u_1, u_2)$.

It is worth pointing out that we are led to exactly *the same mathematical problem* if interested, instead, in another very obvious question, namely, in the analysis of the behavior of the state (q, z) in response to attempts to follow *time-varying tracking signals*, even in the absence of observation errors. Indeed, in that case the d_{ij} 's would be zero, but the different possible tracking functions θ_d and r_d would still give rise to potentially arbitrary inputs u_1 and u_2 . In summary, either of these two basic control questions: sensitivity to measurement error, or analysis of time-varying (instead of merely constant) tracking signals, gives rise to the problem of studying stability of the system with respect to the inputs u .

As in current nonlinear control studies, and specifically as in Marino's work for tracking problems, one may ask then if our system is ISS when u is taken as an input. (Marino required tracking controllers to have an ISS property with respect to disturbances acting on the system. In the special case when the disturbances are matched to the control, this amounts to the problem studied here.) In particular, if the system were to be ISS, then bounded inputs u should result in bounded trajectories (ISS is a stronger property than "bounded-input bounded-state" stability). However, there are bounded inputs which produce *nonlinear resonance* behavior, resulting in unbounded state trajectories, which implies that this system is not ISS. Indeed, the input shown in Figure 4(b) has the property that, for a suitable initial state, the ensuing trajectory is unbounded. Figure 4(c) shows the " r " component of the state of a certain solution which corresponds to this input (see [9] for details on how this input and trajectory were calculated). In conclusion, the tracking system behavior exhibits unstable behavior with respect to measurement disturbances and/or with respect to time-varying reference signals. One might hope, however, given the simplicity and common use of these designs, that *some sort of robustness property is verified* for this system. The study of this question, for this example, led to the work reported in [9]. The answer turns out to have wide applicability. We discovered that the weaker but still very useful property of iISS is always satisfied for the passivity controller in this robotic example, even though ISS is not.

To prove the iISS property, one takes the mechanical energy of the system as a candidate Lyapunov function, and verifies that the following passivity-type estimate holds: $\frac{d}{dt}V(q(t), z(t)) \leq -c_1|z(t)|^2 + c_2|u(t)|^2$ for some sufficiently small number $c_1 > 0$ and some sufficiently large number $c_2 > 0$. Inspection of the equations shows that, when $u \equiv 0$ and $z \equiv 0$, necessarily $q \equiv 0$ as well. Thus, thinking of z as an output, the system is weakly zero-detectable and dissipative; using the theorem cited above, it follows that the system is indeed iISS.

IOSS

The thesis of our student Krichman (completed in December 1999; see also the paper [19]) concerned the following "zero detectability" question: *is it possible to estimate, on the basis of external information provided by past input and output signals, the magnitude of the internal state $x(t)$ at time t ?* State estimation is central to control theory (Kalman filters, observers). By and large, the theory of state estimation is well-understood for linear systems, but it is still poorly developed for more general classes of systems. An outstanding open question is the derivation of useful necessary and sufficient conditions for the existence of observers, i.e., "algorithms" (dynamical systems) which converge to an estimate $\hat{x}(t)$ of the state $x(t)$ of the system of interest, using the information provided by $\{u(s), s \leq t\}$, the set of past input values, and by $\{y(s), s \leq t\}$, the set of past output measurements. In the context of stabilization to an equilibrium, let us say to the zero state $x = 0$ if we are working in an Euclidean space,

a weaker type of estimate is sometimes enough: it may suffice to have a norm-estimate, that is to say, an upper bound $\hat{x}(t)$ on the *magnitude* (norm) $|x(t)|$ of the state $x(t)$. Indeed, it is often the case that norm-estimates suffice for control applications. To be more precise, one wishes that $\hat{x}(t)$ becomes eventually an upper bound on $|x(t)|$ as $t \rightarrow \infty$. We are thus interested in *norm-estimators* which, when driven by the i/o data generated by the system, produce such an upper bound $\hat{x}(t)$. In order to understand the issues that arise, let us start by considering the very special case when the external data (inputs u and outputs y) vanish identically. The obvious estimate (assuming, as we will, that everything is normalized so that the zero state is an equilibrium for the unforced system, and the output is zero when $x = 0$) is $\hat{x}(t) \equiv 0$. However, the only way that this estimate fulfills the goal of upper bounding the norm of the true state as $t \rightarrow \infty$ is if $x(t) \rightarrow 0$. In other words, one obvious necessary property for the possibility of norm-estimation is that the origin must be a globally asymptotically stable state with respect to the “subsystem” consisting of those states for which the input $u \equiv 0$ produces the output $y \equiv 0$. One says in this case that the original system is *zero-detectable*. For *linear systems*, zero detectability is equivalent to detectability, that is to say, the property that if any two trajectories produce the same output, then they approach each other. Zero-detectability is a central property in the general theory of nonlinear stabilization on the basis of output measurements, and is classically found in references such as Isidori’s book). Our work can be seen as a contribution towards the better characterization and understanding of this fundamental concept.

However, zero-detectability by itself is far from being sufficient for our purposes, since it fails to be “well-posed” enough. One easily sees that, at the least, one should ask that, when inputs and outputs are small, states should also be small, and if inputs and outputs converge to zero as $t \rightarrow \infty$, states do too. Moreover, when defining formally the notion of norm-estimator and the natural necessary and sufficient conditions for its existence, other requirements appear: the existence of asymptotic bounds on states, as a function of bounds on input/output data, and the need to describe the “overshoot” (transient behavior) of the state. One way to approach the formal definition, so as to incorporate all the above characteristics in a simple manner, is to dualizing the definition of ISS, arriving at the notion of detectability called *output to state stability* (OSS), or, more generally if there are both inputs and outputs, *input/output to state stability* (IOSS): A system $\dot{x} = f(x, u)$ with measurement (output) map $y = h(x)$ is IOSS if there exist functions $\beta \in \mathcal{KL}$ and $\gamma_1, \gamma_2 \in \mathcal{K}_\infty$ such that the estimate:

$$|x(t)| \leq \beta(|x(0)|, t) + \gamma_1 \left(\|u|_{[0,t]}\| \right) + \gamma_2 \left(\|y|_{[0,t]}\| \right)$$

holds for any initial state $x(0)$ and any input $u(\cdot)$, where $x(\cdot)$ is the ensuing trajectory and $y(t) = h(x(t))$ the respective output function. The terminology IOSS is self-explanatory: formally there is “stability from the i/o data to the state”. It represents a natural combination of the notions of “strong” observability (cf. the PI’s 1989 coprime factorizations paper) and ISS, and was called simply “detectability” in one of our 1989 papers (where it is phrased in input/output, as opposed to state space, terms, and applied to questions of parameterization of controllers) and was called “strong unboundedness observability” by Jiang, Teel, and Praly. (more precisely, this last notion allows also an additive nonnegative constant in the right-hand side of the estimate). In [2], we described relationships between the existence of full state observers and the IOSS property, or more precisely, a property which we called *incremental IOSS*. (The study of incremental IOSS is still in its beginning stages, and represents a major direction for further work.) The use of ISS-like formalism for studying observers, and hence implicitly the IOSS property, has also appeared several times in other authors’ work.

One of the main results of [19] is that a system is IOSS if and only if it admits a norm-estimator. This result is in turn a consequence of a necessary and sufficient characterization of the IOSS property in terms of smooth dissipation functions, namely, there is a proper (radially unbounded) and positive definite smooth function V of states (a “storage function” in the language of dissipative systems) such that a dissipation inequality

$$\frac{d}{dt}V(x(t)) \leq -\sigma_1(|x(t)|) + \sigma_2(|y(t)|) + \sigma_3(|u(t)|) \quad (17)$$

holds along all trajectories, with the functions σ_i of class \mathcal{K}_∞ . This provides an “infinitesimal” description of IOSS, and a norm-observer is easily built from V . Such a characterization in dissipation terms was conjectured for several years, but only recently we were able to obtain the complete solution.

It is worth pointing out that several authors (Lu, Morse, etc) had independently suggested that one should *define* “detectability” in dissipation terms. In some work one finds detectability defined by the requirement that there should exist a differentiable storage function V satisfying our dissipation inequality but with the special choice $\sigma_2(r) := r^2$ (there were no inputs in the class of systems considered there). A variation of this is to weaken the dissipation inequality, to require merely $x \neq 0 \Rightarrow \frac{d}{dt}V(x(t)) < \sigma_2(|y(t)|)$ (again, with no inputs), as done for instance in the definition of detectability given by Morse. Observe that this represents a slight weakening of our property, in so far as there is no “margin” of stability $-\sigma_1(|x(t)|)$. One of our contributions is to show that such alternative definitions (when posed in the right generality) are in fact equivalent to IOSS. (We have recently started work with Morse and Liberzon which ties together the notion of IOSS with the applications in adaptive control which motivated Morse’s work. The paper [44] shows how to formulate the requisite “minimum phase” type of property as an IOSS property for an extended system whose outputs are derivatives of the original output.)

A key preliminary step in the construction of V , just as it was for the analogous result for the ISS property, is the characterization of the IOSS property in robustness terms, by means of a “small gain” argument. The IOSS property is shown to be equivalent to the existence of a “robustness margin” $\rho \in \mathcal{K}_\infty$. This means that every system obtained by closing the loop with a feedback law Δ (even dynamic and/or time-varying) for which $|\Delta(t)| \leq \rho(|x(t)|)$ for all t , is OSS (i.e., is IOSS as a system with no inputs). The core of the paper [19] is, thus, the construction of V for “robustly detectable” (more precisely, “robust IOSS”) systems $\dot{x} = g(x, d)$ which are obtained by substituting $u = d\rho(|x|)$ in the original system, and letting $d = d(\cdot)$ be an arbitrary measurable function taking values in a unit ball. The function V must satisfy a differential inequality of the form $\dot{V}(x(t)) \leq -\sigma_1(|x(t)|) + \sigma_2(|y(t)|)$ along all trajectories, that is to say, the following partial differential inequality:

$$\nabla V(x) \cdot g(x, d) \leq -\sigma_1(|x|) + \sigma_2(|y|),$$

for some functions σ_1 and σ_2 of class \mathcal{K}_∞ . But one last reduction consists of turning this problem into one of building Lyapunov functions for “relatively asymptotically stable” systems. Indeed, one observes that the main property needed for V is that it should *decrease along trajectories as long as $y(t)$ is sufficiently smaller than $x(t)$* . This leads us to the notion of “global asymptotic stability modulo outputs” and its Lyapunov-theoretic characterization.

The construction of V relies upon the solution of an appropriate optimal control problem, for which V is the value function. This problem is obtained by “fuzzifying” the dynamics near the set where $y \ll x$, so as to obtain a problem whose value function is continuous. Several elementary facts about relaxed controls are used in deriving the conclusions. The last major

ingredient is the use of techniques from nonsmooth analysis, and in particular inf-convolutions, in order to obtain a Lipschitz, and from there by a standard regularization argument, a smooth, function V , starting from the continuous V that was obtained from the solution of the optimal control problem.

IOS

The concept of input to *output* stability (IOSS) is the last inhabitant of the “zoo” of ISS-like notions which we consider. Recall from the discussion in the Introduction that we view the components of the triad consisting of ISS, IOS, and IOSS as the generalizations, respectively, of the three basic notions of stability type classical for linear systems: (1) internal stability (behavior of states), (2) detectability (external behavior determines how states evolve), and (3) input/output stability (transfer function has no poles in right-hand plane). They play the exact same role as the respective linear-system notions; the theorem that a system is internally stable if and only if it is i/o stable and detectable generalizes to the fact (basically tautological, once that the definitions have been given) that ISS is the conjunction of IOSS and IOS.

We are currently investigating several concepts which arise as common generalizations of IOSS and IOS, such as the “measurement and inputs to output” stability concept, in which we have two output functions y and w , one of which is seen as a measurement and the other as an objective for stabilization. The corresponding notion, in which $|w(t)|$ is bounded in terms of a \mathcal{KL} function of initial states and time (or more generally, yet another “measure” of states) and nonlinear gains on u and y , can be characterized in dissipation terms, and the proof, on which we are working with a graduate student, should provide a unifying tool in the theoretical development. Further down the line, we expect to formulate results in an even more abstract framework in which inputs and outputs are not distinguished (the “port” model in Willems’ behavioral approach). For now, the concepts are all treated separately. Let us quickly introduce the topic and a motivation from regulation theory, and explain some of the main existing results.

The problem of i/o stability for systems with outputs

$$\dot{x}(t) = f(x(t), u(t)), \quad y(t) = h(x(t)) \quad (18)$$

was studied in [12] and [14], following preliminary results given in [26]. Here $h : \mathbb{R}^n \rightarrow \mathbb{R}^p$, for some p , is locally Lipschitz and satisfies $h(0) = 0$. Roughly, a system (18) is “output stable” if, for any initial state, the output $y(t)$ converges to zero as $t \rightarrow \infty$. Inputs u may influence this stability in different ways; for instance, one may ask that $y(t) \rightarrow 0$ only for those inputs for which $u(t) \rightarrow 0$, or just that y remains bounded whenever u is bounded. Such behavior is of central interest in control theory. As an illustration, we will review below how regulation problems can be cast in these terms, letting $y(t)$ represent a quantity such as a tracking error. Another motivation for studying output stability arises in classical differential equations: “partial” asymptotic stability (in the sense of e.g. Vorotnikov) is nothing but the particular case of our study in which there are no inputs u and the coordinates of y are a subset of the coordinates of x (that is to say, h is a projection on a subspace of the state space \mathbb{R}^n). It is interesting to note that partial stability, or stability of a function of states (output), was already mentioned in Lyapunov’s original work in 1892. Also, as pointed out in [26], the notion of output stability is also related to Lakshmikantham’s “stability with respect to two measures.”

Of course, there are many different ways to make mathematically precise what one means by “ $y(t) \rightarrow 0$ for every initial state” (and, when there are inputs, “provided that $u(t) \rightarrow 0$ ”). These different definitions need not result in equivalent notions; one must decide how

uniform is the rate of convergence of $y(t)$ to zero, and precisely how the magnitudes of inputs and initial states affect convergence. To do this, we follow our general ISS guiding principles of using combinations of \mathcal{KL} -decay and nonlinear gains. There are two technically different, but conceptually equivalent, ways to formulate the property of input/output stability and its variants. One is in purely input/output terms, where one uses past inputs in order to represent initial conditions. Another is in state space terms, where the effect of past inputs is summarized by an initial state. It turns out that the IOS case is substantially more complicated than ISS, in the sense that there are subtle possible differences in definitions. The paper [12] elaborated on these differences, compared the various definitions, and established a theorem on *output redefinition*, which extends one of the main steps in linear regulation theory to general nonlinear systems and provides one of the main technical tools needed for the construction of Lyapunov functions accomplished in [14].

Output regulation problems include the analysis of feedback systems with the following property: for each exogenous signal $d(\cdot)$ (which might represent a disturbance to be rejected, or a signal to be tracked), the output $y(\cdot)$ (respectively, a quantity being stabilized, or the difference between a certain variable in the system and its desired target value) must decay to zero as $t \rightarrow \infty$. Typically, the exogenous signal is unknown but is constrained to lie in a certain prescribed class (for example, the class of all constant signals). Moreover, this class can be characterized through an “exosystem” given by differential equations (for example, the constant signals are precisely the possible solutions of $\dot{d} = 0$, for different initial conditions). Suppose that we already have a closed-loop system exhibiting the desired regulation properties, ignoring the question of how an appropriate feedback system has been designed. Moreover, let us, for this informal discussion, restrict ourselves to linear time-invariant systems (local aspects of the theory can be generalized to certain nonlinear situations employing tools from center manifold theory, see Isidori’s textbook). The object of study becomes

$$\dot{z} = Az + Pw, \quad \dot{w} = Sw, \quad y = Cz + Qw,$$

seen as a system $\dot{x} = f(x)$, $y = h(x)$, where the extended state x consists of z and w ; the z -subsystem incorporates both the state of the system being regulated (the plant) and the state of the controller, and the equation $\dot{w} = Sw$ describes the exosystem that generates the disturbance or tracking signals of interest. (This is a system without inputs; later we explain how inputs may be introduced into the model as well.) As an illustration, take the stabilization of the position y of a second order system $\ddot{y} - y = u + w$ under the action of all possible constant disturbances w . The conventional proportional-integral-derivative (PID) controller uses a feedback law $u(t) = c_1 q(t) + c_2 y(t) + c_3 v(t)$, for appropriate gains c_1, c_2, c_3 , where $q = \int y$ and $v = \dot{y}$. Let us take $c_1 = -1$, $c_2 = c_3 = -2$. If we view the disturbances as produced by the “exosystem” $\dot{w} = 0$, the complete system becomes $\dot{q} = y$, $\dot{y} = v$, $\dot{v} = -q - y - 2v + w$, $\dot{w} = 0$ with output y . Writing the plant/controller state z as $\text{col}(q, y, v)$, we can write this in matrix form for with $S = Q = 0$ and appropriate A, C, P . In linear regulator theory, the routine way to verify that the regulation objective has been met is as follows. Suppose that the matrix A is Hurwitz and that there is some matrix Π such that the following two identities (“Francis’ equations”) are satisfied:

$$\Pi S = A\Pi + P, \quad 0 = C\Pi + Q.$$

(The existence of Π is necessary as well as sufficient for regulation, provided that the problem is appropriately posed.) Consider the new variable $\hat{y} := z - \Pi w$. The first identity for Π allows decoupling \hat{y} from w , leading to $\dot{\hat{y}} = A\hat{y}$. Since A is a Hurwitz matrix, one concludes that

$\hat{y}(t) \rightarrow 0$ for all initial conditions. As the second identity for Π gives that $y(t) = C\hat{y}(t)$, one has the desired conclusion that $y(t) \rightarrow 0$.

Let us now express this convergence in a much more informative form. For that purpose, we introduce the map $\hat{h} : x = (z, w) \mapsto |z - \Pi w| = |\hat{y}|$. We also denote, for ease of future reference, $\chi(r) := r/|C|$ and $\beta(r, t) = r |e^{tA}|/|C|$, using $|\cdot|$ to denote Euclidean norm of vectors and also the corresponding induced matrix norm. So, $y = C\hat{y}$ gives

$$\chi(|y|) = \chi(|h(x)|) \leq \hat{h}(x) = |\hat{y}| \quad (19)$$

for all x , and we also have $|y(t)| \leq \beta(|\hat{y}(0)|, t)$ and in particular

$$|y(t)| \leq \beta(|x(0)|, t), \quad \forall t \geq 0 \quad (20)$$

along all solutions. This estimate quantifies the rate of decrease of y to zero, and its overshoot, in terms of the initial state of the system. For the auxiliary variable \hat{y} , we have in addition the following “stability” property:

$$|\hat{y}(t)| \leq \sigma(|\hat{y}(0)|), \quad \forall t \geq 0 \quad (21)$$

where $\sigma(r) := r \sup_{t \geq 0} |e^{tA}|$.

The use of \hat{y} (or equivalently, finding a solution Π for the above matrix identities) is a key step in the analysis of regulation problems. Note the fundamental contrast between the behaviors of \hat{y} and y : because of (21), a zero initial value $\hat{y}(0)$ implies $\hat{y} \equiv 0$, which in regulation problems corresponds to the fact that the initial state of the “internal model” of the exosignal matches exactly the one for the exosignal; on the other hand, for the output y , typically an error signal, it may very well happen that $y(0) = 0$ but $y(t)$ is not identically zero. The fact about \hat{y} which allows deriving (20) is that \hat{y} dominates the original output y , in the sense of (19). One of the main results of [12] provides an extension to very general nonlinear systems of the technique of output redefinition.

To illustrate again with the PID example: one finds that $\Pi = \text{col}(1, 0, 0)$ is the unique solution of the required equations, and the change of variables consists of replacing q by $q - w$, the difference between the internal model of the disturbance and the disturbance itself, and $\hat{h}(q, y, v, w) = (|q - w|^2 + |y|^2 + |v|^2)^{1/2}$. For instance, with $x(0) = y(0) = v(0) = 0$ and $w(0) = 1$ we obtain the output $y(t) = \frac{1}{2}t^2e^{-t}$. Notice that this output has $y(0) = 0$ but is not identically zero, which is consistent with an estimate (20). On the other hand, the dominating output $\hat{y} = (q - w, y, v)$ cannot exhibit such overshoot.

The discussion of regulation problems was for systems $\dot{x} = f(x)$ which are subject to no external inputs. This was done in order to simplify the presentation and because classically one does not consider external inputs. In general, however, one should study the effect on the feedback system of perturbations which were not exactly represented by the exosystem model. A special case would be, for instance, that in which the exosignals are not exactly modeled as produced by an exosystem, but have the form $w + u$, where w is produced by an exosystem. Then one may ask if the feedback design is robust, in the sense that “small” u implies a “small” asymptotic (steady-state) error for y , or that $u(t) \rightarrow 0$ implies $y(t) \rightarrow 0$. Experience with the notion of ISS then suggests that one should replace (20) by an estimate as follows:

$$(\text{IOS}) \quad |y(t)| \leq \beta(|x(0)|, t) + \gamma(\|u\|).$$

By this we mean that for some functions γ of class \mathcal{K} and β of class \mathcal{KL} which depend only on the system being studied, and for each initial state and control, such an estimate holds for

the ensuing output. We suppose as a standing hypothesis that the system is forward-complete, that is to say, solutions exist (and are unique) for $t \geq 0$, for any initial condition and any locally essentially bounded input u . The reason that inputs u are not usually incorporated into the regulation problem statement is probably due to the fact that for *linear* systems it makes no difference: it is easy to see, from the variation of parameters formula, that IOS holds for all inputs if and only if it holds for the special case $u \equiv 0$. Property (21) generalizes when there are inputs to the following “output Lagrange stability” property:

$$(OL) \quad |y(t)| \leq \sigma_1(|y(0)|) + \sigma_2(\|u\|).$$

As mentioned earlier when discussing the classical tools of regulation theory, one of our main results is this: if a system satisfies IOS, then we can always find *another output*, let us call it \hat{y} , which dominates y , in the sense of (19), and for which the estimate OL holds in addition to IOS. As $\dot{\hat{y}} = A\hat{y}$ and A is a Hurwitz matrix, in the case of linear regulator theory the redefined output \hat{y} satisfies a stronger decay condition, which in the input case leads naturally to an estimate as follows:

$$(SIIOS) \quad |y(t)| \leq \beta(|y(0)|, t) + \gamma(\|u\|)$$

(we write y instead of \hat{y} because we wish to define these notions for arbitrary systems). For linear systems, the conjunction of IOS and OL is equivalent to SIIOS. Remarkably, this equivalence breaks down for general nonlinear systems. Finally, as for ISS, there are close relationships between output stability with respect to inputs, and robustness of stability under output feedback. This suggests the study of yet another property, which is obtained by a “small gain” argument from IOS: there must exist some $\chi \in \mathcal{K}_\infty$ so that

$$(ROS) \quad |y(t)| \leq \beta(|x(0)|, t) \quad \text{if} \quad |u(t)| \leq \chi(|y(t)|) \quad \forall t.$$

For linear systems, this property is equivalent to IOS, because applied when $u \equiv 0$ it coincides with IOS. One the main contributions of [12] was the construction of a counterexample to show that this equivalence also fails to generalize to nonlinear systems. In summary, we showed in [12] that precisely these implications hold:

$$SIIOS \Rightarrow OL \ \& \ IOS \Rightarrow IOS \Rightarrow ROS$$

and that under output redefinition the two middle properties coincide.

The paper [14] provided necessary and sufficient dissipation characterizations for each of the properties defined above, for systems that are uniformly bounded-input bounded-state. Among many other results, an equivalence is established between IOS and the existence of some smooth function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ such that, for some $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$, $\alpha_1(|h(x)|) \leq V(x) \leq \alpha_2(|x|)$ for all states x and, for some $\chi \in \mathcal{K}$, $DV(x)f(x, u) < 0$ whenever $V(x) \geq \chi(|u|)$ (for all states x and control values u).

*List of Papers (Journals, Book Chapters, Conference Proceedings) on this Project,
Submitted or Appeared During the Grant Period*

1. F.H. Clarke, Yu.S. Ledyaev, E.D. Sontag, A.I. Subbotin, "Asymptotic controllability implies feedback stabilization," *IEEE Trans. Autom. Control* **42**(1997): 1394-1407.
2. E.D. Sontag, Y. Wang, "Output-to-state stability and detectability of nonlinear systems," *Systems and Control Letters* **29**(1997): 279-290.
3. E.D. Sontag, F.R. Wirth, "Remarks on universal nonsingular controls for discrete-time systems," *Systems and Control Letters* **33**(1998): 81-88.
4. E.D. Sontag, "Comments on integral variants of ISS," *Systems and Control Letters* **34**(1998): 93-100.
5. Yu.S. Ledyaev, E.D. Sontag, "A Lyapunov characterization of robust stabilization," *J. Nonlinear Analysis* **37**(1999): 813-840.
6. W. Desch, H. Logemann and E.P. Ryan, E.D. Sontag, "Meagre functions and asymptotic behaviour of dynamical systems," *J. Nonlinear Analysis*, to appear.
7. D. Nešić, E.D. Sontag, "Input-to-state stabilization of linear systems with positive outputs," *Systems and Control Letters* **35**(1998): 245-255.
8. D. Angeli, E.D. Sontag, "Forward completeness, unboundedness observability, and their Lyapunov characterizations," *Systems and Control Letters* **38**(1999): 209-217.
9. D. Angeli, E.D. Sontag, Y. Wang, "A characterization of integral input to state stability," *IEEE Trans. Autom. Control* **45**(2000): 1082-1097.
10. D. Nešić, E.D. Sontag, A.R. Teel, "Formulas relating \mathcal{KL} stability estimates of discrete-time and sampled-data nonlinear systems," *Systems and Control Letters* **38** (1999): 49-60.
11. F. Albertini, E.D. Sontag, "Continuous control-Lyapunov functions for asymptotically controllable time-varying systems," *Int. J. Control* **72**(1999): 1630-1641.
12. E.D. Sontag, Y. Wang, "Notions of input to output stability," *Systems and Control Letters* **38**(1999): 235-248.
13. E.D. Sontag, "Clocks and insensitivity to small measurement errors," *Control, Optimisation and Calculus of Variations* **4**(1999): 537-557.
14. E.D. Sontag, Y. Wang, "Lyapunov characterizations of input to output stability," *SIAM J. Control and Opt.* **39**(2001): 226-249.
15. L. Grune, E.D. Sontag, F.R. Wirth, "Asymptotic stability equals exponential stability, and ISS equals finite energy gain - if you twist your eyes," *Systems and Control Letters* **38**(1999): 127-134.
16. D. Angeli, E.D. Sontag, Y. Wang, "Further equivalences and semiglobal versions of integral input to state stability," *Dynamics and Control* **10**(2000): 127-149.
17. M. Malisoff, E.D. Sontag, "Universal formulas for feedback stabilization with respect to Minkowski balls," *Systems and Control Letters* **40**(2000): 247-60.

18. L. Rosier, E.D. Sontag, "Remarks regarding the gap between continuous, Lipschitz, and differentiable storage functions for dissipation inequalities appearing in H_∞ control," *Systems and Control Letters* 41(2000): 237-249.
19. M. Krichman, E.D. Sontag, Y. Wang, "Input-output-to-state stability," *SIAM J. Control and Optimization*, to appear.
20. D. Liberzon, A.S. Morse, E.D. Sontag, "An ISS-like variant of the minimum-phase property," *IEEE Trans. Autom. Control*, to appear.
21. E.D. Sontag, "A general approach to path planning for systems without drift," in *Essays on Mathematical Robotics* (J. Baillicul, S.S. Sastry, and H.J. Sussmann, eds.), Springer-Verlag, 1998, pp. 151-168.
22. E.D. Sontag, "Control-Lyapunov functions," in *Open Problems in Mathematical Systems and Control Theory*, (V.D. Blondel, E. Sontag, M. Vidyasagar, and J.C. Willems, eds.) Springer Verlag, London, 1999, pp. 211-216.
23. E.D. Sontag, "Nonlinear Feedback Stabilization Revisited," in *Dynamical Systems, Control, Coding, Computer Vision* (Proc. Math. Theory of Networks and Systems, Padova, July 1998; G. Picci and D.S. Gilliam, eds.), Birkhäuser Verlag, Basel, 1999, pp. 223-262.
24. E.D. Sontag, "Stability and stabilization: Discontinuities and the effect of disturbances," in *Nonlinear Analysis, Differential Equations, and Control* (Proc. NATO Advanced Study Institute, Montreal, Jul/Aug 1998; F.H. Clarke and R.J. Stern, eds.), Kluwer, Dordrecht, 1999, pp. 551-598.
25. E.D. Sontag, "The ISS philosophy as a unifying framework for stability-like behavior," in *Nonlinear Control in the Year 2000* (Volume 2) (Lecture Notes in Control and Information Sciences, A. Isidori, F. Lamnabhi-Lagarigue, and W. Respondek, eds.), Springer-Verlag, Berlin, 2000, pp. 443-468.
26. E.D. Sontag, Y. Wang, "A notion of input to output stability," *Proc. European Control Conf.*, Brussels, July 1997, Paper WE-E A2, CD-ROM file ECC958.pdf, 6 pages.
27. F. Albertini, E.D. Sontag, "Control-Lyapunov functions for time-varying set stabilization," *Proc. European Control Conf.*, Brussels, July 1997, Paper WE-E A5, CD-ROM file ECC515.pdf, 6 pages.
28. E.D. Sontag, "Recent results on discontinuous stabilization and control-Lyapunov functions," in *Proc. Workshop on Control of Nonlinear and Uncertain Systems*, London, Feb. 1998.
29. E.D. Sontag, "Notions of integral input-to-state stability," *Proc. American Control Conf.*, Philadelphia, June 1998, pp. 3215-3219
30. Yu. Ledyayev, E.D. Sontag, "Stabilization under measurement noise: Lyapunov characterization," *Proc. American Control Conf.*, Philadelphia, June 1998, pp. 1658-1662
31. D. Angeli, E.D. Sontag, Y. Wang, "A remark on integral input to state stability," in *Proc. IEEE Conf. Decision and Control, Tampa, Dec. 1998*, IEEE Publications, 1998, pp. 2491-2496.

32. D. Nešić, E.D. Sontag, "Output stabilization of nonlinear systems: Linear systems with positive outputs as a case study," in *Proc. IEEE Conf. Decision and Control, Tampa, Dec. 1998*, IEEE Publications, 1998, pp. 885-890.
33. M. Krichman, E.D. Sontag, "A version of a converse Lyapunov theorem for input-output to state stability," in *Proc. IEEE Conf. Decision and Control, Tampa, Dec. 1998*, IEEE Publications, 1998, pp. 4121-4126.
34. M. Malisoff, E.D. Sontag, "Universal formulas for CLF's with respect to Minkowski balls," in *Proc. American Control Conf.*, San Diego, June 1999, pp. 3033-3037.
35. D. Liberzon, E.D. Sontag, Y. Wang, "On integral-input-to-state stabilization," in *Proc. American Control Conf.*, San Diego, June 1999, pp. 1598-1602.
36. D. Nešić, E.D. Sontag, A.R. Teel, "On stability and input-to-state stability \mathcal{KL} estimates of discrete-time and sampled-data nonlinear systems," in *Proc. American Control Conf.*, San Diego, June 1999, pp. 3990-3994.
37. Z-P. Jiang, E.D. Sontag, Y. Wang, "Input-to-state stability for discrete-time nonlinear systems," in *Proc. 14th IFAC World Congress (Beijing)*, Vol E, pp. 277-282, 1999.
38. E.D. Sontag, "Feedback insensitive to small measurement errors," in *Proc. IEEE Conf. Decision and Control, Phoenix, Dec. 1999*, IEEE Publications, 1999, pp. 2661-2666.
39. B. Ingalls, E.D. Sontag, Y. Wang, "Remarks on input to output stability," in *Proc. IEEE Conf. Decision and Control, Phoenix, Dec. 1999*, IEEE Publications, 1999, pp. 1226-1231.
40. M. Krichman, E.D. Sontag, Y. Wang, "Lyapunov characterizations of input-output-to-state stability," in *Proc. IEEE Conf. Decision and Control, Phoenix, Dec. 1999*, IEEE Publications, 1999, pp. 2070-2075.
41. D. Angeli, E.D. Sontag, "Characterizations of forward completeness," in *Proc. IEEE Conf. Decision and Control, Phoenix, Dec. 1999*, IEEE Publications, 1999, pp. 2551-2556.
42. L. Grune, E.D. Sontag, F.R. Wirth, "On the equivalence between asymptotic and exponential stability, and between ISS and finite H_∞ gain," in *Proc. IEEE Conf. Decision and Control, Phoenix, Dec. 1999*, IEEE Publications, 1999, pp. 1220-1225.
43. M. Chyba, N.E. Leonard, E.D. Sontag, "Time-optimal control for underwater vehicles," in *Lagrangian and Hamiltonian Methods for Nonlinear Control* (N.E. Leonard, R. Ortega, eds), Pergamon Press, Oxford, 2000, pp. 117-122.
44. D. Liberzon, A.S. Morse, E.D. Sontag, "A new definition of the minimum-phase property for nonlinear systems, with an application to adaptive control," in *Proc. IEEE Conf. Decision and Control, Sydney, Dec. 2000*, IEEE Publications, 2000, pp. 2106-2111.
45. (with B. Ingalls and Y. Wang) "Generalizations of asymptotic gain characterizations of ISS to input-to-output stability," *Proc. American Control Conference*, June 2001, to appear.
46. (with D. Angeli and Y. Wang) "A Note on Input-to-State Stability with Input Derivatives," *Proc. Nonlinear Control System Design Symposium*, July 2001, to appear.